

FEATURE GUIDE

Double-Take[®] for Linux

TABLE OF CONTENTS

Introduction3
Product Description	3
Features3
Intelligent Compression	3
Email Alerts	3
Asynchronous, Block or Byte-level, Continuous Replication	3
Hardware Independence	4
Bandwidth Throttling	4
Simple Management Console	4
Uses Existing Networks	4
Task Command Processing	4
Unlimited Distance Replication	4
Many-to-One Failover	5
SNMP Counters and Traps	5
Automated Failover	5
Open-file Mirroring and Replication	5
Flow Control (Unlimited Disk Queuing)	5
Serverless Backup Support	5
Block Checksum Re-Mirror	5
Replication Scheduling	6
Automatic Re-Mirror	6
File/Directory Selection	6
Support for Different File Types	6
Verification	6
Full Command-Line Control	6
Broad Distribution Support	6
About Double-Take® Software7

Introduction

Product Description

Double-Take for Linux 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 uses patented replication and failover technology that continuously captures block or 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.

Asynchronous, Block or 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. Block or 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 infrastructure 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.

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.

Simple Management Console

The 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.

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.

Automated Failover

Double-Take can provide high-availability failover of servers to ensure users remain online in case of a failure. Double-Take has no distance limitations, 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 server, 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.

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.

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 on the Linux operating system.

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.

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.

Broad Distribution Support

Double-Take offers support for Red Hat Enterprise Linux 4 and 5, Novell SUSE Linux Enterprise 10, and CentOS 4 and 5. Support for Linux includes both 32-bit and 64-bit support with the ability to mirror and replicate between both architectures and other supported distributions (RH and CentOS). The Ext2 and Ext3 file systems are supported for all distributions in addition to ReiserFS being supported for Novell SUSE Linux Enterprise.

About Double-Take® Software

Headquartered in Southborough, Massachusetts, Double-Take® Software (Nasdaq: DBTK) is a leading provider of affordable software for recoverability, including continuous data replication, application availability and system state protection. Double-Take Software products and services enable customers to protect and recover business-critical data and applications such as Microsoft Exchange, SQL, and SharePoint in both physical and virtual environments. With its unparalleled partner programs, technical support, and professional services, Double-Take Software is the solution of choice for more than ten thousand customers worldwide, from SMEs to the Fortune 500. Information about Double-Take Software's products and services can be found at www.doubletake.com.

Manage your subscription to eNews. Visit: www.doubletake.com/subscribe

Double-Take Software Headquarters

257 Turnpike Road
Southborough, MA 01772
Phone: +1-800-964-0185 or +1-508-229-8483
Fax: +1-508-229-0866

Double-Take Software Sales

8470 Allison Pointe Blvd. Suite 300
Indianapolis, IN 46250
Tel: +1-888-674-9495 or +1-317-598-0185
Fax: +1-317-572-1867

Double-Take Software EMEA

Tel: +33 147 77 05 00
saleseu@doubletake.com

For Sales in UK and Ireland

Tel: + 44 (0) 1905 745711
saleseu@doubletake.com

For Sales in Germany

Tel: +49 69 6897776-66
salesgermany@doubletake.com

For Sales in Benelux

Tel: + 32 2 209 16 16
salesblx@doubletake.com

For Sales in Latin America

+1.317.572.1859 or +1.888.674.9495
saleslatam@doubletake.com

For Sales in Asia Pacific, Japan, India and Australia/New Zealand, excluding China

+65 6866 3308
salesap@doubletake.com

For Sales China

Tel: +86 21 6109 5747
salesap@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 Double-Take Software, Inc. The information in this document is subject to change without notice. Although we try to provide quality information, Double-Take Software 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.



Get the standard today: www.doubletake.com or 888-674-9495

© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.