Feature Guide

Double-Take Availability for Linux

Double-Take® Availability 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 Availability for Linux 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, you can from your target servers in minutes, if not seconds. What's more, Double-Take Availability for Linux 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.

Intelligent Compression Double-Take Availability for Linux is the only host-based replication solution that provides multilevel intelligent compression. Four data compression options allow custom configurations and settings can be individually configured for each replication set defined, allowing even further customization by server, data, and/or network.

Choose 'no compression' 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 additional CPU and RAM use 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.

As certain data types can actually increase in size when compressed, resulting in more data being transmitted rather than less, Double-Take Availability for Linux compares the size of the compressed data to that of the uncompressed and transmits the smaller of the two, ensuring that compression does not increase the load on the network. Double-Take Availability for Linux will intelligently compress only the data from which a benefit will result.

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 Availability for Linux 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 committed to the target disk and the corresponding acknowledgement has been sent back to the production server.

Double-Take Availability for Linux 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.

Hardware Independence Whereas hardware-based synchronous replication is often proprietary for each storage vendor and provides no cross-vendor replication, Double-Take Availability for Linux works with standard Direct attached or SAN attached hardware with no affinity to a specific hardware manufacturer or model. This independence allows Double-Take Availability for Linux 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, 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 Availability for Linux, and define the data to be replicated.

Bandwidth Throttling User-defined controls limit the amount of the available network bandwidth used for data replication, allowing replication to occur real-time without impacting users on the same network. And, since Double-Take Availability for Linux 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 Availability for Linux allows the administrator to define the actual type of network connection and then define the portion that can be used for replication. Double-Take Availability for Linux also allows you to limit network use during busy work hours and increase or remove limits during non-peak hours.

Replication Console The Recovery Console automatically discovers servers running Double-Take Availability for Linux and displays them in a single window. Managing the entire deployment can be performed from this single console and any transfers, reports, failures, and/or failbacks can be viewed from this single point. Double-Take Availability for Linux also includes a Full-server Backup Job Creation Wizard that simplifies setting up a full-server backup to an image server. Administrators can now export existing replication sets as templates to any folder where the Replication Console is installed. These templates can then be imported and used to simplify setting up new connections.

Uses Existing Networks Double-Take Availability for Linux does not require its own private network for replicating data. Most often, existing networks are more than sufficient, allowing you to implement low-cost data replication and protection solutions and reducing the overall cost-of-ownership by keeping the initial investment lower and removing any restrictions and additional costs for future network changes or purchases. Double-Take Availability for Linux can also be run on its own private network should you want to isolate replication traffic.

Task Command Processing Double-Take Availability for Linux 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, 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 Availability for Linux uses standard IP networks, there aren't the typical distance limitations as with some synchronous and array-based products. 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 Availability for Linux can be configured such that a single target server is used for numerous source servers, reducing the total cost of ownership and facilitating centralized backup.

Email Alerts For optimal data protection and availability, Double-Take Availability for Linux 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 Availability for Linux 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 Availability for Linux, 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 via email, administrators do not have to continually monitor the Recovery Console to ensure everything is operating smoothly.

Double-Take Availability for Linux Feature Guide

SNMP Counters and Traps Double-Take Availability for Linux will forward both replication statistics and events via SNMP allowing simplified management by administrators.

Automated Failover Double-Take Availability for Linux can provide high-availability failover of servers to ensure users remain online in case of a failure. There are 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.

Open-file Mirroring and Replication Double-Take Availability for Linux doesn't require that applications be restarted each time additional files and/or directories are configured for replication. Users, customers and applications remain online. Double-Take Availability for Linux 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 Availability for Linux 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 Availability for Linux will continue to filter all file changes and buffer these changes while transmitting to the target. Queuing ensures that all transactions are replicated to the target without loss of data. With the bandwidth throttling and scheduling features, users can configure 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 applied once the backup is complete.

Serverless Backup Support 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 Availability for Linux can facilitate data restoration from the target back to the original source or to an alternate location. Through the Recovery Console, users can easily restore data from the replicated disk back to the production disk once the failure is corrected, ensuring that you recover from the time the failure occurred, not from when the last backup was taken. Unlike other solutions that make the users remember which files came from what location, the restore process automatically reverses the direction of your original replication job.

Full-Server Protection Capabilities Double-Take Availability for Linux provides the ability to replicate changes to both user and application data as well as operating system components, application packages and other aspects of the protected server's system state. This continuously updated disk-based copy of the server can be used to reduce the time needed to recover to new hardware in the event of a system failure.

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 Availability for Linux 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.

Replication Scheduling Double-Take Availability for Linux provides the flexibility to schedule when replication occurs. Double-Take Availability for Linux 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 Availability for Linux 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 Availability for Linux 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.

Double-Take Availability for Linux Feature Guide

File/Directory Selection You can select full file systems, directories or files for replication, providing 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 Availability for Linux has been successfully tested against different file types on the Linux operating system.

Verification Although Double-Take Availability for Linux 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 Availability for Linux allows all GUI functions to be controlled from the command line, either via scripts or individual commands, giving you greater flexibility and automation capabilities.

TCPBuffer Auto-tuning
Double-Take Availability for Linux allows the O/S to dynamically auto-tune the TCP Buffer Size eliminating the need to have the user fine tune for a specific network and allowing for the most efficient use of the available bandwidth.

Broad Distribution Support Double-Take Availability for Linux offers support for Novell SUSE Linux Enterprise 9,10 and 11 . 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 (Red Hat and CentOS). The Ext2 and Ext3 file systems are supported for all distributions in addition to ReiserFS being supported for Novell SUSE Linux Enterprise.

Automatic Re-Mirror Should a problem or even scheduled maintenance require that the connection between the source and target be broken, Double-Take Availability for Linux 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.

Brisbane Level 2, Suite 4, 67 Astor Terrace, Spring Hills, QLD 4000, Australia
Tel: +61 7 31771000 Fax: +61 2 9383 8658

Sydney Suite 501, Level 5 55 Grafton St, Bondi Junction, NSW 2022
Tel: +61 2 8644 3502 Fax: +61 2 9383 8658

Melbourne Suite 2323, Level 23, HWT Tower, 40 City Road, SouthGate, VIC 3006
Tel: +61 3 9090 9134 Fax: +61 2 9383 8658

business continuity

visit www.doubletake.com.au

redhat. READY

Ready FOR SUSE® LINUX ENTERPRISE Novell.