

MIGRATING TO VIRTUAL MACHINES WITH THE DOUBLE-TAKE® VIRTUAL RECOVERY ASSISTANT

TECHNICAL WHITEPAPER

Double-Take Software, Inc.

Published: June 2008

Abstract

Though server virtualization technology has changed the IT landscape for the better in many ways, it has also presented several new challenges. One of the chief concerns that IT decision makers face is how to most effectively migrate between disparate server and storage infrastructures without halting production applications or requiring extensive manual intervention. Using a combination of real-time data replication and full-server failover technologies for workload migration can be an effective way to address this challenge.

Introduction

Though server virtualization technology has changed the IT landscape for the better in many ways, it has also presented several new challenges along the way. One of the chief concerns that IT decision makers face is how to most effectively migrate between disparate server and storage infrastructures without halting production applications or requiring extensive manual intervention. These problems are compounded exponentially for migration of applications over Wide Area Networks (WANs). WANs dramatically reduce the amount of available bandwidth when compared to Local Area Network (LAN) migrations and latency adds a high level of uncertainty for data integrity. The Double-Take Virtual Recovery Assistant was designed specifically for high latency environments with patented algorithms that ensure data integrity regardless of available bandwidth or the latency of the network connection. It also includes features for migration and failover of applications so that your business doesn't have to wait hours to continue working. Rather, downtime is limited to the few moments that it takes to start the applications again. Users can continue working while you're migrating workloads and you don't have to worry about production server impact because it is easy to control the allocation resources available to Double-Take replication and failover. After the migration of workloads into new virtual machines (VMs) is complete you can also continue using Double-Take or the Double-Take Virtual Recovery Assistant for high availability and off-site disaster recovery projects for both virtual and physical servers. The primary benefit to this approach is the lack of having to re-learn a new toolset or spend large amounts of money and time on tools that require a lot of training.

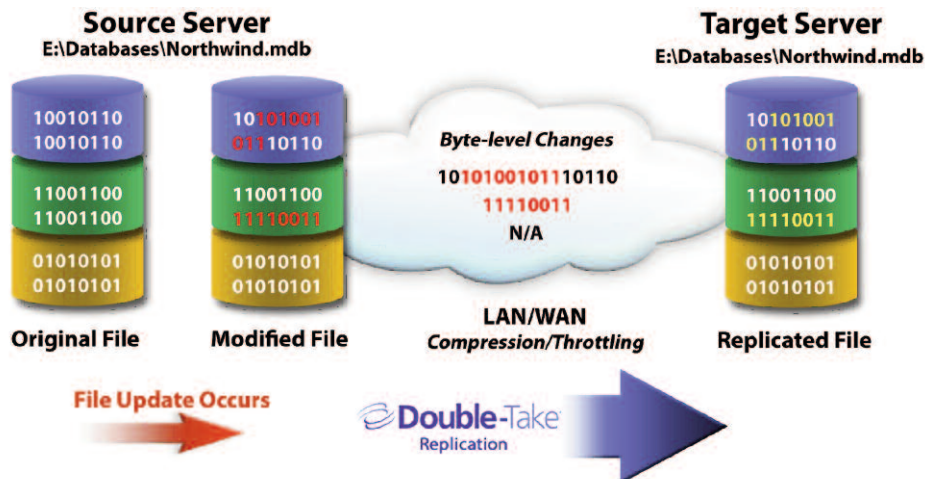
What Double-Take for Windows Provides

Double-Take replication technology reduces the window for data loss, is more efficient than tape, and less costly than tape or other replication solutions. Because Double-Take replicates data in near real-time, an up-to-the-second copy of data is always available on a second server. Double-Take also uses unique compression and bandwidth throttling features that allow an administrator to control the amount and timing of bandwidth used for backup. Double-Take can utilize existing infrastructure and is hardware agnostic; there's no large investment or hidden costs standing between you and better backups.

Double-Take allows for the complete recovery of any Microsoft Windows application and the data associated with it. Many Double-Take customers replicate their data across continents and oceans to provide nearly instantaneous recovery regardless of the cause of the outage. It combines continuous real-time replication and automatic failover capabilities for disaster recovery, high availability, and centralized backup of applications. The real-time protection of Double-Take provides a Recovery Point Objective (RPO) usually measured in seconds.

Today, Double-Take is the most relied-upon solution for real-time replication of critical data and automated failover for application availability. Double-Take is Microsoft® Windows® 2000 and 2003 certified at all levels, one of the few replication products to have achieved this level of certification. It delivers protection that is better or comparable to many hardware based solutions, but costs tens of thousands less.

Double-Take replicates changes to files at the byte-level from any Windows Server to any other Windows Server across any IP-based network. It installs on each server and uses a file system filter driver to determine what changes are occurring to files and it replicates those file operations to another server and applies them to a secondary copy of the data. All changes are sent and applied in the exact order that they occurred on the production system, guaranteeing a crash-consistent copy of data on the secondary system.



Double-Take Benefits

- **Real-Time Protection** - Replicates continuously at the byte level over any shared or private IP network, ensuring that changed data is protected and can be quickly restored at all times
- **Application Agnostic** - Works with existing hardware to protect applications such as Microsoft Exchange, Microsoft SQL Server, and SharePoint® – any application that runs on Windows Server.
- **Easily Installed and Maintained** - Allows companies of any size looking for data protection solutions to install and maintain Double-Take.
- **Cost Effective** - Provides strong data protection at low cost with an accelerated return on investment - paying for itself usually within months.

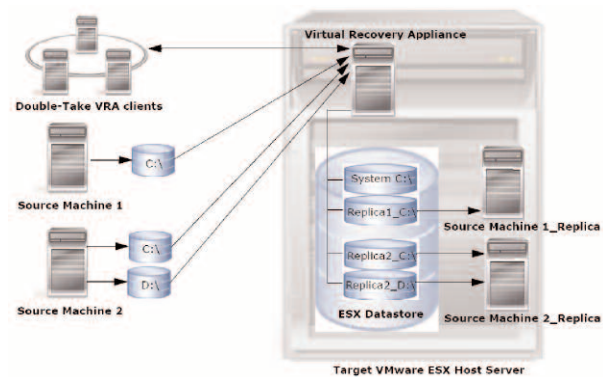
How the Virtual Recovery Assistant Works

The Double-Take Virtual Recovery Assistant builds on the award winning replication capabilities of Double-Take for Windows which has been tested and proven reliable in enterprise data centers for well over a decade. This core replication platform provides the data integrity functionality necessary to guarantee that production applications never need to go off-line to migrate from their production servers to their new virtualized home. By leveraging Double-Take data replication and workload failover technologies, the Double-Take Virtual Recovery Assistant provides all of the functionality necessary to migrate data, applications and the operating system in real-time. Then it can provision the new virtual machine and bring it online within minutes to complete the migration while completely mitigating some of the risks associated with other migration tools.

The Double-Take Virtual Recovery Assistant Architecture

The Double-Take Virtual Recovery Assistant architecture is based on a client/server model that includes three primary components:

- **Virtual Recovery Appliance:** A virtual machine which runs on a target VMware ESX Server host and performs the migration functionality including managing virtual disks, push-installing Double-Take for Windows to source servers and performing migration job management.
- **Target VMWare ESX host:** A host server running VMware ESX Server which runs one or more Virtual Recovery Appliances and provides the hosting platform for migrated servers.
- **Management Clients:** These typically run on administrator desktops and manage the Virtual Recovery Assistant Appliances and their relationship with the VMware ESX host server(s) as well as the migration, protection and the recovery processes of moving source servers into newly created replica VMs.



These three systems work together to migrate source servers to VMs in real-time without interruption of the production server application(s).

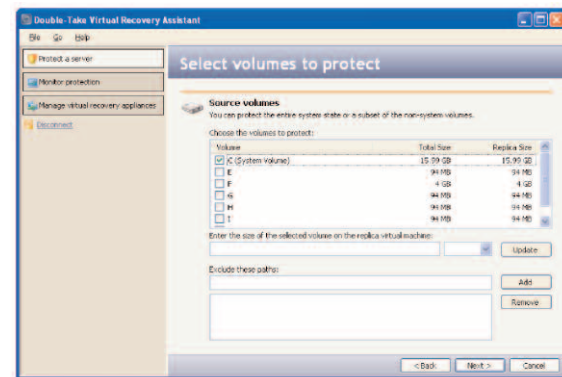
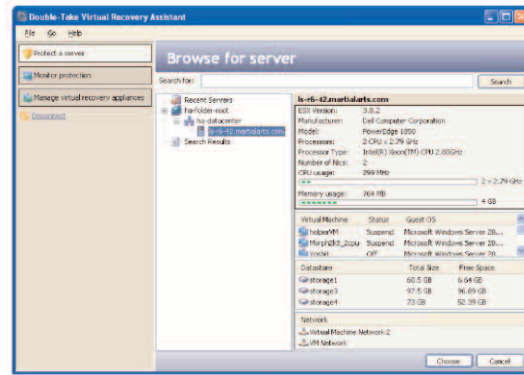
- **Source servers:** Production systems that are to be migrated to a target VMware ESX host server. These can be either physical servers or virtualized servers on non-VMware platforms.
- **Replica Virtual Machines (VMs):** Virtual machines (VMs) into which source servers are migrated. They remain powered off until failover is triggered to complete the migration process and bring them online.

Implementing Virtual Recovery Assistant

The Double-Take Virtual Recovery Assistant provides real-time migration for any Windows 2000 or Windows 2003 server, regardless of whether it is implemented as a physical or virtual server. The Virtual Recovery Assistant also works with all popular virtual server platforms including Microsoft Virtual Server, Microsoft Hyper-V, Xen and Virtual Iron. This provides the ultimate in flexibility without sacrificing production application downtime because of disparate server or storage platforms.

Installing the Virtual Recovery Assistant begins by installing the management client on a desktop to manage the migration process. Then you will provision a Virtual Recovery Appliance on the Target VMware ESX server that you will use to host your migrated servers. This can be a VMware ESX server that already has existing workloads running on it or a newly provisioned VMware ESX host created specifically to host migrated workloads.

Next, start the Virtual Recovery Assistant console and connect to a VMware VirtualCenter server with credentials appropriate for viewing, creating and modifying new or existing virtual machines within the datacenter. Then select a source server that you would like to migrate. Double-Take is required on the source server because it is used for data movement as part of the migration process. If Double-Take is not already installed on the machine, then the Virtual Recovery Assistant can push-install it on the administrator's behalf. The first step in configuring a migration job is to select the volumes that you would like to protect or migrate. Double-Take provides unique features to ignore all non-recoverable I/O on the production system so that you don't have to waste time and bandwidth because of large non-recoverable files like Pagefile.sys or temporary database files for applications like SQL Server and Exchange. It also safely ignores volume and disk management activity like disk defragmentation. Other migration products force you to migrate and protect all the data on the production system even though it may have no value in the protection or migration process. Double-Take has the ability to selectively migrate only the changes that are required. This will save you a lot of time to baseline your machines and dramatically reduces the cost, especially for migration projects across Wide Area Networks.



After selecting the source server and the volumes which will be replicated, you must specify which target VMWare ESX host to use for running the newly created virtual machine. The console allows you to see the CPU, memory, storage and network usage on each target VMWare ESX host server. In combination with the Virtual Recovery Assistant's ability to auto-discover ESX Servers in the environment, you can easily locate and select host servers which have enough capacity to handle the migrated workload from the source server.

Once you have located a suitable VMWare ESX host server, you can then select an appropriate target datastore which will hold the virtual disk files for the newly created VM. These virtual disk files will eventually contain all of the data from the production source server being migrated.

The next step gives you the opportunity to select the VM configuration options such as:

- **Replica Name:** A name to be given to the newly created virtual machine. This is used for identifying and managing the VM within the administrative tools provided by VMware.
- **Replica Network:** This allows you to specify which virtual network that you want the replica VM to attach to after migration completes.
- **Processors and Memory:** You can use the same amount of memory and number of processors as the production source server or change these parameters as needed to maximize and optimize the performance of the replica VM on the target ESX Server host.

You also have the opportunity to select additional options such as compression which can dramatically reduce the total network bandwidth requirements. Since this feature is based on the Double-Take replication engine, it requires very few resources and delivers as high as 80% compression of real-time transactional data like SQL Server, Oracle and Exchange.

Next, select the failover options to be used when migrating from the existing source servers to the new replica Virtual Machine. Some of these options include specifying the production IP Address to monitor for failure and whether or not to failover to the replica Virtual Machine automatically or manually. While automated failover might seem advantageous at first glance, it is rarely used in practice. This allows the administrator to have complete control over the migration process. If you have a large server that might require considerable time to initially synchronize to the replica Virtual Machine, you do not want to force migration if a network outage were to occur during the business day. Using the manual failover feature during migration will allow you to choose precisely when the migration process completes.

Finally, you are presented with a protection summary that lists all the actions that the migration job will perform and allows you to review before accepting and initiating the process.

Managing Virtual Recovery Assistant

Once you have created a protection job for your migration project, then you're able to manage the active migration project through the Double-Take Virtual Recovery Assistant. It lists all of the active migration or protection jobs and shows their status. Once their status is in an "Idle" state, you can perform failover at any time.



Many users find that their protection jobs may complete during the business day and choose not to failover immediately to reduce the impact to production users and to maintain compliance with any change control management parameters in place.

Regardless of when you choose to complete the migration process, Double-Take for Windows will continue to actively replicate only the byte-level changes to your production systems. If SQL Server changes 1 byte in its database files, for example, Double-Take will send only a single byte across the network. Other migration products on the market send the entire 64 Kilobyte disk block because they can't see the logical transaction and merely send the disk-level changes that occurred as a result of the operation.

Completing Migration

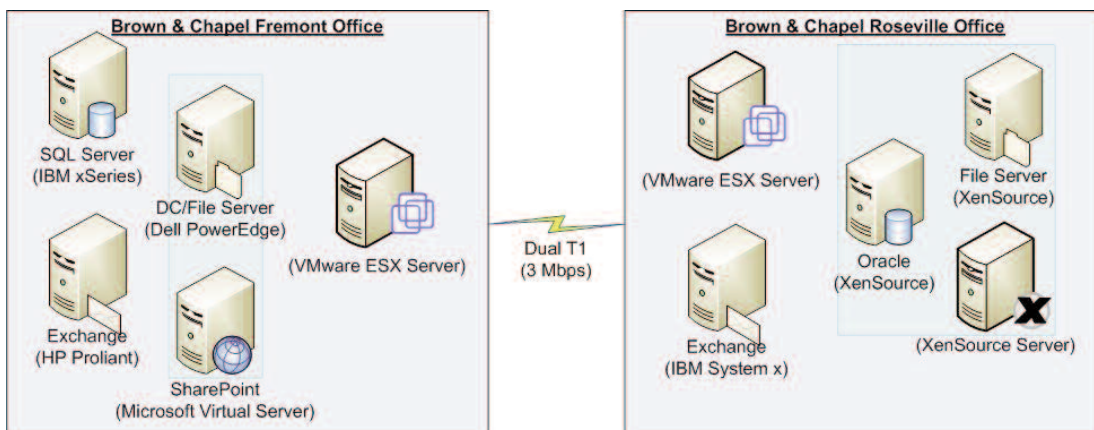
Once you are ready to complete the migration process from your Source Server to the Virtual Replica Machine, then you can take the production server off the network. Next, pick the appropriate Migration Job and select "Failover". Now the Virtual Recovery Assistant will dynamically perform the following tasks:

- VMware virtual drivers are injected into the Replica Virtual Machine
- The virtual disks are dismounted from the Virtual Recovery Appliance
- The virtual disks are mounted to the Replica Virtual Machine
- The replica Virtual Machine is powered on

When the replica Virtual Machine completes booting, you can log into it and see that your data, applications and operating system are in a good and consistent state including any changes to the registry, event logs and application-related data. All aspects of the new virtual machine will appear as they did on the production source server. The entire actual migration process completes in the time that it takes to boot the new VM. Administrators can say "goodbye" to lost weekends and wasted overtime.

A Migration Case Study Example

Let's take an opportunity to apply the information above in two use case scenarios. We'll look at how Brown & Chapel, a fictitious regional engineering firm with 130 employees located in two offices in Fremont, CA and Roseville, CA. Brown & Chapel would like to achieve two primary objectives to reduce their IT support expenditures and prepare for easier future growth. First, they want to migrate from their existing disparate server and storage infrastructure to VMWare ESX Server to consolidate servers and take advantage of a new server's capacity. Their next objective is to provide off-site data protection and recoverability to protect their data and applications at both of their offices in the event of a site outage, which will allow their billable engineers to continue working from home if the office isn't available.



B&C's primary location has a mix of hardware that they purchased over the years and have implemented a virtual machine for SharePoint Services using Microsoft Virtual Server 2005 and their existing file server. They would like to migrate all of their servers to the VMware ESX platform, but don't want to endure extensive interruption of their applications due to large billable projects that their engineers are actively working on.

Their Roseville office was an acquired company with a mix of platforms, but had already begun the process of moving their servers to the Xen virtualization platform when acquired. They want to standardize on a single platform so those servers will also be migrated to a VMware ESX Server that they have purchased for that office. After server migrations have completed, B&C engineering recognizes the need to get their data out of each office for DR purposes in a more timely and consistent manner than their previous strategy of administrators taking tape backups home at night.

Phase 1: Migrating Disparate Systems to VMware

The Brown & Chapel IT team decides to come in an hour earlier in the morning before most of their coworkers begin work for the day and implement the Double-Take Virtual Recovery Assistant to begin the migration process. This gives them the ability to reboot the machines after installing Double-Take and still allow their engineers to continue working on their critical project late into the evening without interruption.

The VMware ESX Servers are running on new dual processor, quad core CPU's and have 32 GB of RAM and plenty of disk storage for growth. Brown & Chapel's IT staff begins migration of the production servers in the Fremont office. They install the Virtual Recovery Assistant Management Console on their administrators' laptops and then create a new virtual machine that will run the Virtual Recovery Appliance. Next they select each production server in turn and decide how they want to deploy them in their new virtual homes. The Exchange server has been running sluggish for the past few months because the production server is running out of processing resources and could really use two dedicated cores. They also know that their SQL Server could certainly use more memory than the 2 GB that it was previously available, so they can increase that to 4 GB. That will still leave them with more than enough resources to run their production servers in Fremont as well as for hosting Roseville's applications during an outage and room for growth.

Once the migration parameters are specified then the IT staff begins push-installing Double-Take to their production servers and starts the baseline mirroring process. Since they are using Double-Take compression features, the synchronization can happen faster than hardware-speed and they maximize their throughput of all the server migrations on the gigabit LAN. Meanwhile, the users have continued working while the synchronization process takes place. At one point, the Exchange server begins to experience some network slowdown when the engineers start arriving and begin emailing large technical drawings between their offices. The Exchange server has only a single NIC so the administrators decide that they will throttle the Double-Take network traffic to maintain acceptable performance levels. Alternatively, they could have pushed the migration traffic to a secondary NIC and avoided the public user interface entirely. They use the Double-Take Management Console to reduce the maximum bandwidth availability to 25 gigabytes per hour and give Exchange Server the additional network resources that it needs during this peak time.

The other great benefit of Double-Take is its extensively tuned algorithms for replicating real-time changes in parallel with the baseline synchronization mirror. Data is protected as soon as new email or transactions are posted to their Exchange and Database servers, which greatly reduces the total amount of time required to perform the entire migration process. Some of the other solutions that B&C considered queued up those changes on the production machine, more than quadrupled the stress on the production server's I/O sub-system, and in the end sent the same data more than once. They also didn't have to endure downtime and wait while the entire system was copied to the new virtual machines using other migration tools. Using Double-Take Virtual Recovery Assistant for server migration cut their migration time by more than half.

While data synchronization at the Fremont location is happening, the IT staff kicks off the migration process at the Roseville office from remotely over their 10 Mbps connection. They monitor the throughput and response time of their applications and once they know that everything is migrating at a comfortable level then they're free to continue working on other daily tasks at hand. Just to be on the safe side, they use the native Windows Performance Monitor's alerting features and set some thresholds for their production application performance. They also decide to forward any Double-Take event log errors using the native Windows Task Scheduler so that they can further adjust throughput if necessary and to notify them when a migration completes. They won't complete the migration process until they get an open maintenance window later this evening.

Later in the afternoon, the IT staff logs into the Virtual Recovery Assistant Console and clicks the Failover button for the Exchange server and watches as it is gracefully migrated into the new virtual machine. After a couple of minutes, the Outlook client on the Administrator's laptop reconnects and continues to check for email. One of the administrators sends a test message from his personal account and it arrives shortly after without having made any changes. The other servers migrate successfully in turn.

Phase 2: Offsite Data Protection and Recoverability

The day after the server migrations, the Brown & Chapel IT staff decide that they should begin implementing an offsite data protection strategy since they now had all of their eggs in one basket on a single VMware ESX Server in each location.

Again they turned to the Double-Take Virtual Recovery Assistant and configured new jobs that sent their virtualized servers in Fremont to Roseville. The process of performing the initial mirror would take much longer than their migration had taken the previous day because they had only 10 Mbps of shared bandwidth available. While they wouldn't see the performance of their LAN migration, they knew that Double-Take wouldn't impact production and they took advantage of its bandwidth shaping technologies to control how much bandwidth was used for synchronization and replicate. They began mirroring the servers, limiting the total throughput to half of the 10Mbps, which would allow a slow but steady trickle of the baseline server mirroring to Roseville without inhibiting the normal production usage of their bandwidth-hungry end-users when transferring technical drawings. After hours they would allow Double-Take to scale up and use the entire 10Mbps connection before throttling it back again at 8:00 am.

Two weeks later, B&C experienced an outage on their power grid and the power company said that it would take a couple of days to repair. They had already completed the large project, but now their normal business was going to suffer for a couple of days while the power was restored. IT contacts the Roseville IT staff and asked them to log into Double-Take Virtual Recovery Assistant and trigger a fail-over. This brought their applications back on-line in Roseville within a few minutes. They told the engineers with laptops that they could use their VPN software to get back into the network and continue working from home. For those without laptops, they used Outlook Web Access and set up Windows Terminal Services so that they could keep using email and gain remote access to their data. This allowed B&C to continue working effectively even though their production servers were no longer available.

Once power was restored later that week and they could reopen the office, they began to resynchronize the new data in Roseville back to their Fremont office while the users kept using it. Their Exchange and SQL Server databases were rather large, but Double-Take's patented algorithms were smart enough to skip all of the data that was already on the original Fremont server and just resynchronize the binary differences. This meant that their re-mirroring time wasn't as long as the original baseline. After the re-mirror completed, then they clicked the failover button and resumed the application back in Fremont.

Thus Brown & Chapel Engineering was able to achieve both of their goals for performing server virtual migration and off-site data protection and recoverability. The best part of all was that they didn't have to spend a lot of money and learn how to use multiple tools to achieve this solution. Everything that they needed was built into Double-Take Virtual Recovery Assistant.

Summary

One of the major barriers to moving existing physical servers into virtual machines is the downtime associated with the migration. Existing tools lack real-time data replication capabilities and require the physical server to be taken offline for several hours while the data is moved to the new virtual machine. Because Double-Take and the Virtual Recovery Assistant can replicate changes to data in real-time, end users can continue to access production applications right up until the moment at which the workload is migrated.

The Double-Take Virtual Recovery Assistant reduces the complex, often manual, steps required to ready the new virtualized environment for migration. It does the work of provisioning the target virtual machine for the user- eliminating the need to setup the new virtual machine, install an operating system, patches or applications. For workloads which need more room to grow the Virtual Recovery Assistant also allows users to easily increase parameters such as disk capacity, allocated memory, and the number of processors available to the application once it is running in a virtual machine.

Most organizations understand the benefits of moving to a virtualized environment but can't afford the downtime associated with migrations or don't have the time or resources required to move their workloads from the physical world to a virtual one. This new feature essentially eliminates the costly downtime associated with migrations.

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
Tel: +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: +1-317-598-0187

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, Australia and New Zealand

Tel: +65 6866 3308
Sales: salesap@doubletake.com

For sales in China

Tel: +86 21 6109 5747
Sales: salesap@doubletake.com



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

