



RecoverTrac™ v2.0
Technical Whitepaper
Service-Oriented Disaster Recovery

FalconStor®

Service-Oriented Disaster Recovery **RecoverTrac™ v2.0 Technical Whitepaper**

FalconStor Software, Inc.
2 Huntington Quadrangle, Suite 2S01
Melville, NY 11747
Phone: 631-777-5188
Fax: 631-501-7633
Website: www.falconstor.com


Copyright © 2011 FalconStor Software. All Rights Reserved.

FalconStor Software, FalconStor, RecoverTrac, MicroScan, TimeMark, DiskSafe, and their respective logos are trademarks or registered trademarks of FalconStor Software, Inc. in the United States and other countries.

Windows is a registered trademark of Microsoft Corporation.


All other brand and product names are trademarks or registered trademarks of their respective owners.

FalconStor Software reserves the right to make changes in the information contained in this publication without prior notice. The reader should in all cases consult FalconStor to determine whether any such changes have been made.



Contents

Introduction	1
Standard DR Solutions	1
Addressing Challenges	2
Overview and Benefits	3
The Two Faces of RecoverTrac – Solo or Integrated	3
How Does RecoverTrac Work?	3
Production Site	3
DR Site.....	4
Key Features.....	5
Configuring RecoverTrac	7
Use Cases	9
Local Bare Metal Recovery (P2P, V2V, or P2V).....	9
Remote Test/Development Team Data Refresh	9
Periodic DR Drill/Rehearsal	10
Automated Remote Site Failover/Failback for DR (P2P, V2V, or P2V/V2P).....	10
Site Migration for Workload Balancing and Workload Distribution	11
Service Providers Offering DR as a Service (DRaaS)	11
VMware Site Recovery Manager Failover/Failback	11
iSCSI Boot Recovery	11
Conclusion	12



Introduction

Protecting your data means nothing if you can't recover it. One of the greatest data center challenges today is ensuring a smooth recovery of operations after downtime. Downtime can be caused by data loss or corruption, equipment failure, or a complete site outage after a loss of power or a natural disaster. In particular, resuming operations at a disaster recovery (DR) site, whether planned (such as a scheduled site migration) or unplanned (such as an accidental event) requires careful preparation. The planning will take months, but the execution of the plan needs to occur within minutes. During these precious minutes, all of the teams involved will be put under pressure to carry on their recovery procedures in a coordinated fashion. Anything could go wrong during the dozens, if not hundreds, of steps performed by the application, hardware, network, and storage teams. A human error, a process flaw, a routing issue, pretty much anything could delay the site recovery. For many, this herculean effort is, simply put, risky and unpredictable.

FalconStor believes that DR automation should not be this complex, and has a great deal of expertise and experience in this sector. As a result, FalconStor has engineered RecoverTrac™, a universal DR orchestrator that automates complex DR tasks, bringing service-oriented recovery to both physical and virtual server infrastructures. This technical white paper discusses the concerns associated with DR, and explains in detail how RecoverTrac technology addresses a full range of DR challenges.

Standard DR Solutions

Various solutions from various vendors have been designed to orchestrate the DR workflow. VMware, for example, has designed its Site Recovery Manager product to enable simple, one-button-recovery and site failover execution. The VMware Site Recovery Manager leverages array-based replication to send protected data to the DR site, but it does not offer application-level consistency of the protected data. The result is crash-consistent recovery. The data for most of the supported arrays must be provisioned as primary storage. To take advantage of VMware Site Recovery Manager, applications must also be running inside virtual machines (VMs) hosted on the VMware ESX Server hypervisor. Other applications, either running on physical servers or on other hypervisor platforms, cannot be included as part of such a recovery plan. Site failback, the ability to move a failed-over workload from the DR data center back to the production data center, is not yet possible. Citrix has an equivalent solution called Site Recovery for XenServer, with similar capabilities.

Microsoft has taken a slightly different approach to DR by allowing supported applications to be clustered via Microsoft Windows Server 2008 R2 Failover Clustering configured as multi-site clusters, which is an adaptation of Microsoft Cluster Service (MSCS). Failover clustering leverages third-party arrays, storage virtualization gateways, or host-based replication. Microsoft environments supported by MSCS, including Microsoft SQL Server and Hyper-V, can have nodes located in different subnets, allowing for stretched-cluster workloads across data centers. Microsoft Windows applications that support MSCS can be protected in this manner, as well as MSCS clustered Microsoft Hyper-V hosted VMs, which can run Microsoft Windows or Linux workloads. Using Microsoft Hyper-V with Microsoft Windows Multi-site Clusters enables VM-based applications that are not MSCS aware to get high availability functionality supported by two different data center sites. Replication over a WAN in this scenario will cause the workload to pause and automatically restart within a few minutes as part of its VM cluster support.

But what about environments where the production data center has a heterogeneous mix of physical and virtual machines? How do you orchestrate a successful recovery? And what if the DR site and the production data center operate on different compute and storage platforms?

Unlike these solutions, which have challenges in dealing with heterogeneous environments, RecoverTrac can handle physical-to-physical (P2P), virtual-to-virtual (V2V), and physical-to-virtual (P2V) automated recovery jobs for any certified hypervisor or physical platform, in any type of data center, even with multiple network segments. Recovery jobs can include local data recovery, such as bare metal recovery


jobs, as well as remote data recovery, with both site failover and site failback orchestration. Whether your data sits on a SAN or on local disks, the FalconStor® Continuous Data Protector (CDP) solution can protect, replicate, and recover your data, even in heterogeneous SAN environments. As a key component of FalconStor CDP, RecoverTrac orchestrates and automates the entire recovery process, including changing the IP addresses of recovered hosts to match the new location and network subnet and performing the conversion process to allow a physical workload to boot and run inside a VM.

Addressing Challenges

To summarize the typical and common challenges faced by data center and IT administrators when it comes to data recovery, and in particular, DR planning and execution, here's a short list of questions, which the latest version of RecoverTrac was designed to address. Do any of these sound familiar?

- How do I protect physical machines? Do I have to recover to the identical physical machine at the DR site? Can I use a VM as the DR standby machine? Can this be fully automated?
- The concept of VMware Site Recovery Manager is great, but I can't implement it because I don't have the same storage array on both sites / my storage array doesn't have a Storage Replication Adapter / my array is not certified by VMware. Can your solution provide complete recovery coverage? During a site outage requiring a site failover, I want to make sure the data that I replicated from the primary data center is recovered, without any data loss or corruption, in the quickest manner possible. Most solutions I have looked at only guarantee crash-consistent data. Can I make sure my tier-one database is mounted cleanly without wasting time repairing inconsistencies?
- During a site outage, I may want to only temporarily move my protected workloads, physical machines, and VMs to the DR site, using RecoverTrac to automate the process. Can RecoverTrac also automate the entire failback process back to the production site? If one of my production servers is a physical system, and I use RecoverTrac to move it to the DR site as a VM, can I move it back and convert it to run again on the original physical machine after the failback?
- How are VMware ESX Server, Citrix XenServer, and Microsoft Hyper-V VM recoveries supported? How is site failover and failback performed when using heterogeneous hypervisor environments?
- During a site failover, I want to be able to recover to a particular point in time, rather than using the latest replicated state. Is this type of recovery possible?
- I want to rehearse/test a DR site failover scenario involving physical and/or virtual machines. It usually takes days or weeks to schedule this kind of test, and many hours of labor to get the test accomplished and validated, so seldom perform DR rehearsals. Can you really automate a DR test? Is it that simple?
- From the production data center, our engineering team wants our Microsoft SQL servers to be replicated and tested in one DR site, while the accounting department wants the file/print servers to be replicated and restored in another DR site. Can RecoverTrac handle one-to-many, many-to-one, and many-to-many site configurations, on top of the expected one-to-one site mapping?
- Is RecoverTrac easy to manage? Are the recovery tasks managed like backup jobs? Can I get different views based on site location, recovery jobs, or server/hosts? Can I identify each recovery task by job ID? Can I save and restore my RecoverTrac configuration, including the recovery jobs, to a different RecoverTrac server?

The answer to all of these questions is a simple, resounding yes. RecoverTrac was designed to address all of the challenges and concerns that IT professionals face in terms of DR and DR planning.



Overview and Benefits

The Two Faces of RecoverTrac – Solo or Integrated

RecoverTrac not only can become the sole orchestrator for all of your DR site failover/failback needs, but it can integrate with existing environments that already have a DR automation solution in place, such as VMware Site Recovery Manager. As mentioned, VMware Site Recovery Manager was conceived from the get-go to provide automated site failover for VMware VMs. In many data centers, those VMs represent 80-85% of all of the application workloads that need DR protection. The remaining 15-20% are physical servers, which for one reason or another, cannot be virtualized.

Think about all of the HP/UX, Solaris Sparc, and IBM AIX systems that VMware technology cannot virtualize due to the processor architectural difference. What about those robust, highly-demanding X86 Windows or Linux servers that the application administrative team simply refuses to virtualize? RecoverTrac can complement VMware Site Recovery Manager and take care of these remaining and highly important servers. In fact, for the x86 physical servers, RecoverTrac can offer three recovery choices for DR site failover:

1. If VMware Site Recovery Manager is already installed, RecoverTrac can go as far as pre-staging a P2V conversion of the physical server at the production site to create a standby VM instance of it that the VMware solution can actually fail over as part of its recovery plans. This extends VMware Site Recovery Manager's capabilities beyond protecting and failing over only VM workloads.
2. RecoverTrac can help automate the site failback process. RecoverTrac can stage the entire site failover, and keep the physical workload as a physical workload at the DR site.
3. RecoverTrac can stage the entire site failover, including performing a P2V conversion at the destination DR site, to turn the physical workload into a VM. The VM can be VMware, Microsoft Hyper-V, or Citrix XenServer.

How Does RecoverTrac Work?

RecoverTrac works in conjunction with FalconStor CDP. FalconStor CDP is a data protection solution that provides unified backup and DR capabilities, providing fast recovery to any known good point in time, and RecoverTrac is the empowering technology behind automated recovery. Thus, FalconStor CDP must already be protecting application servers and replicating data across data centers before RecoverTrac can add any value to the recovery process.

RecoverTrac can perform both local recovery (bare metal recovery) and remote recovery of servers. For remote recovery, RecoverTrac can handle many-to-many site mappings. For example, two data centers, one in Boston and one in Chicago, could be replicating to a joint remote site in Miami. Or in a different configuration, a single production data center in New York may need to segregate and split its recovery groups into two separate remote DR sites in different cities based on load capacity.

For the sake of simplicity, we will assume two sites: a production site and a DR site.

Production Site

At the production site, an application workload is running with either Fibre Channel (FC) or iSCSI connectivity. To protect the data and ensure application-consistent recovery points, a FalconStor CDP protection package is installed (set it and forget it). For Microsoft Windows systems, this package includes DiskSafe™ technology from FalconStor, as well as application-specific snapshot agents. (A complete list of supported applications and databases can be found at <http://www.falconstor.com/certification-matrix/applications-and-databases>).

A FalconStor CDP appliance or FalconStor CDP Virtual Appliance is installed locally with enough storage capacity to track changes for snapshots and replication purposes. This local appliance backs up data from the protection application workload using changed-block tracking technology (MicroScan™), with a flexible scheduling option in either continuous or periodic mode. (Think of your traditional backup model, except that the schedule can be as frequent as every 10 or 15 minutes, instead of restricted to once a day). The resulting backup points are called TimeMark® snapshots, and they are point-in-time quiesced with full application consistency. The snapshots for each application workload can be replicated remotely to a FalconStor CDP appliance or appliances at one or several DR sites.

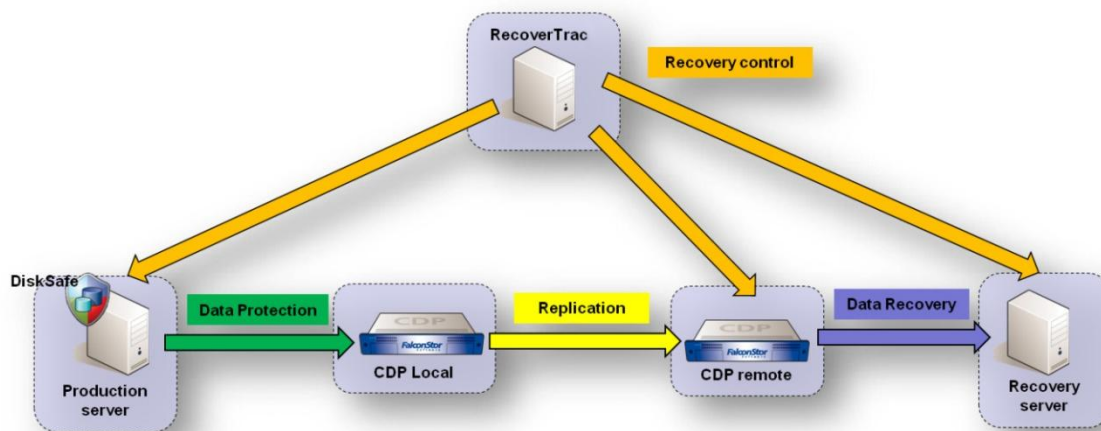
For local recovery, an optional RecoverTrac server can be installed to orchestrate recovery of workloads by communicating with the FalconStor CDP appliance and automating the local data restore process.

DR Site

At the recovery site, a physical server is installed, with the same or similar hardware specifications (same make and model) as the physical server in the production site. Alternatively, a supported hypervisor server capable of running VMs with acceptable performance may be present.

Another FalconStor CDP appliance or FalconStor CDP Virtual Appliance is used as a target appliance. Protected data from the application workload at the production site is replicated continuously or periodically from the local appliance to the target appliance, with quiescent point-in-time snapshots for fast, application-consistent data recovery.

Another RecoverTrac server is installed, as the one optionally installed at the production site could become unavailable during site outages. This RecoverTrac server will orchestrate DR recovery for all of the protected workloads coming from the production site(s).



The two FalconStor CDP appliances, local and target, are basically the backup units. The FalconStor CDP software protection package, installed in each protected workload, performs the scheduling for point-in-time, transactionally consistent snapshots, as well as delta data movement, with changed-block tracking enabled from the primary disk to the local appliance.

FalconStor CDP has its own dedicated management interface, which allows you to perform the initial configuration for replication schedule and site pairing. Beyond the initial deployment, you can use this console if you wish to perform your own manual recovery of data by manually creating a mountable snapshot, assigning that synthetic restored disk as a drive letter or mount point to your designated server, and copying the data back. However, you don't need to perform manual recovery, since RecoverTrac automates every step, while providing many other compelling features.

Once the two FalconStor CDP appliances have been configured to replicate between one another, and the application workload has been protected with DiskSafe™ and the optional snapshot agents, you can then configure the RecoverTrac servers to authenticate to both appliances, such that either can perform recovery tasks and jobs for both sites.


The concept behind RecoverTrac is simple: RecoverTrac carefully de-couples the machines, which define the hardware (virtual or physical), and the application workloads or host images, which define the software (operating system and application data). RecoverTrac can orchestrate the move of the application workload (host image) from a “protected machine” to a “recovery machine” by defining a pair of machines, each mapped to the same host image, and then specifying which machine is the protected machine and which one is the recovery machine.

Key Features

- Ability to integrate with both FalconStor CDP for out-of-band data protection and FalconStor® Network Storage Server (NSS) for in-band storage virtualization.
- FalconStor CDP and FalconStor NSS replication engines leverage MicroScan, a changed-block tracking engine, which allows for sector-level (512-bytes) delta data replication, as well as compression and encryption for the most secure and efficient block-based replication available. This considerably reduces the recovery point objectives (RPO), as replication can be performed more frequently and using less network bandwidth.
- Application workloads (host images) are recovered with full transactional consistency, not just crash consistency, thanks to FalconStor Snapshot Agent integration with your application’s database applications. This allows RecoverTrac to perform recovery with the smallest possible Recovery Time Objective (RTO), and with zero data loss.
- Heterogeneous storage support, heterogeneous remote site storage replication.
- Heterogeneous hypervisor support, V2V conversion across supported hypervisors, and across local and remote sites.
- P2V site failover support, from supported physical server models to supported hypervisor platforms.
- V2P site failback support, from identical paired physical and virtual machines used during the initial site failover.
- P2P site failover and site failback support, from supported physical server models to matching physical server models.
- V2V site failover and site failback support, in any combination of hypervisor platforms.
- Supports FC Boot from SAN. If you are performing a P2V site failover, and the protected machine boots the host image using FC SAN Boot via a FalconStor NSS FC SAN resource, RecoverTrac will allow a V2P site failback to the original protected machine, and continue to use Boot from SAN upon failback. If performing a P2P site failover, and both the protected machine and the recovery machine are using identical physical servers, then the recovery machine will also be able to boot the recovered OS disk over FC SAN Boot.
- Event/Audit Log Views: All recovery operations are logged for a complete audit trail.
- Re-home, which is the ability to adapt the application workload (host image) to changes in the environment after a site failover/migration. For example, a recovery machine may no longer have the same IP address as the original protected machine because it is now connected to a different VLAN with a different gateway. Also, if the application workload was protected using FalconStor NSS and related snapshot agents from FalconStor, the new recovery machine must get its snapshot notifications from the FalconStor NSS appliance used for DR. Finally, if the recovery machine is connected to FalconStor NSS data disks over iSCSI, it also must be reconfigured to reconnect to the

iSCSI target at the target appliance located at the DR site (not the iSCSI target at the production site). RecoverTrac automates this process to ensure a smooth migration/transition to the recovery site, and to prepare for a smooth site failback.

- Clone Mode support, to branch off a machine independently off its source. Useful for short- or long-term testing/development projects, or for cloud service deployments from templates.
- Test Mode support, to rehearse DR recovery plans.
- Ability to change IP address of each host image during a recovery job execution, to accommodate the network scheme of the DR site. This is needed if your DR site is not in the same stretched VLAN as the production site.
- Ability to change back the IP address of each host image upon site failback, as RecoverTrac retains all of the machine information for all of the sites in its database.
- Ability to execute recovery jobs with any available TimeMark snapshot. You do not need to select the very last replication point for the site failover. This is especially important to prevent rolling disasters, such as the primary site data being corrupted after an identified time.
- Ability to schedule a periodic refresh of the recovery jobs to automatically allow the testing/development team to access the latest set of test data after a certain designated period of time, or to automatically pre-stage the earlier steps of a DR orchestration (to accelerate the site failover process).



Configuring RecoverTrac

RecoverTrac features an interface that is friendly and intuitive. This allows the configuration workflow to run in the following manner:

1. Define all of the sites, and specify which site is the local site relative to where this specific instance of RecoverTrac is installed. For example, if this instance of RecoverTrac is installed at the DR site, then the DR site is considered the local site for this RecoverTrac server.

Note: You can define multiple sites, and even multiple RecoverTrac servers per sites, but each RecoverTrac server will manage its own recovery jobs.

2. In the storage servers section, define each site's FalconStor CDP appliance(s), including network info, login credentials, and site location.
3. Optionally, if using VMs in either site, or planning to use RecoverTrac to perform a P2V recovery, you need to add and define access to each site's hypervisor servers (such as VMware ESX Server and Microsoft Hyper-V), along with the respective management centers (such as VMware vCenter). This information will help RecoverTrac automate many of the virtualization steps, such as creating a new VM, or powering a VM on or off. This is done under the 'hypervisor servers and centers' section.
4. Under Hosts and Clusters, you will add all of the machines (virtual or physical) that will be linked to a given Host image (identified by OS hostname).

- a. The first machine you add for a given Host image will identify the production physical or virtual machine that hosts the protected application, which is identified as the Protected Machine.

Certain information must be provided:

- OS hostname, aka Host image (OS name, management IP address, OS type, site location. If the machine is a VM, you must also indicate the host hypervisor server.)
- OS administrator username and password.
- The SAN client object name used by this machine in the FalconStor CDP appliance protecting it.
- The FalconStor CDP disk devices attached directly to the machine. The Boot device must be identified as well, along with its disk geometry (for P2V conversions).
- Network info for remote shutdown.

- b. The remaining machines, which we refer to as recovery machines for the same Host image can now be added to the list, following the same wizard-driven format. You will need to specify the same OS hostname as the protected production machine. However, the site location will differ if the recovery machines are located at the DR site. The machine's IP address(es) may also differ, in the case of a remote recovery, due to network segmentation.


Note: Once a machine is mapped to a host image, it will no longer be available for mapping to other host images. So the same host image can be mapped to multiple machines (and in multiple sites), but each machine can only be mapped to a single host image. Multiple machines for the same host image can be a combination of physical and virtual machines. When performing P2P recovery, all of the machines must be physical servers with nearly identical hardware specifications (same make, same model, and same generation hardware).

5. Finally, define your recovery jobs. A recovery job can contain one or many machines to be recovered, and you can even induce a power on delay between each machine to ensure that any dependencies between applications and servers are respected. In the recovery job creation process, you provide:
 - a. The host images being restored (one or more), and the desired recovery site for each.

- b. For each Host image, you will need to map the protected site's protected disk devices to the recovery site's replica recovery disk devices.
- c. The desired Recovery Mode (Test, Recovery, or Clone).

Additional policies to fine-tune the automation. For example:

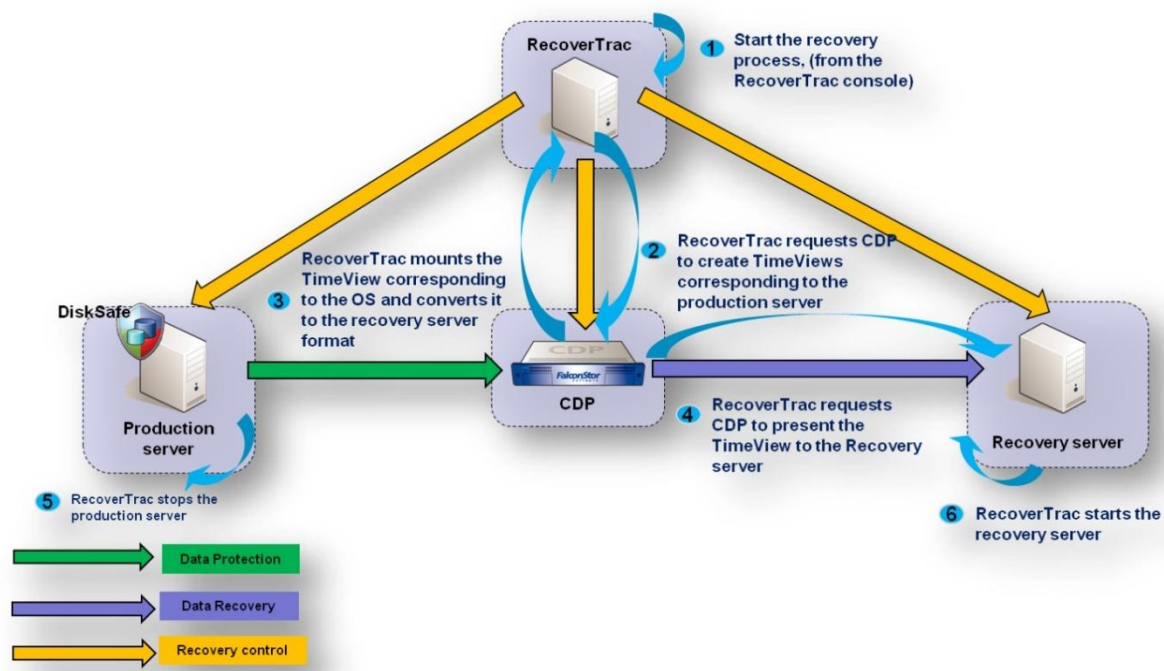
- Do you want to replicate the final delta changes from the protected machine (between the two sites) before performing a site failover?
- Do you want to power on the recovery machines?
- In the case of a virtual recovery machine, do you want to automatically install/update VM tools?
- If multiple machines are being recovered, do you want to stop on the first error encountered, or continue to recover the rest of the machines?



Use Cases

Local Bare Metal Recovery (P2P, V2V, or P2V)

RecoverTrac allows you to use a recovery machine located at the production site to run a recovery job. This restores the host image, and brings the application workload previously running on the protected machine back online, at the recovery machine, without any manual steps. If the protected machine is a VM, then the recovery machine also needs to be a VM. If the protected machine is a supported physical server, then you can define a recovery job and select either a VM or a supported physical server as the recovery machine. (The physical server must be of similar hardware, make, and model.)



Remote Test/Development Team Data Refresh

Testing and development teams in other regions of the world often need a copy of the production data to perform data mining, testing, and analysis. RecoverTrac allows you to schedule special recovery jobs, which will bring recovery machines online at the testing/development remote data center. These recovery machines which will contain a carbon copy of the protected machine from a specific point in time. This TimeMark snapshot can be selected from all of the snapshot states that were replicated, or you can specify a schedule to perform a dataset refresh, if needed. The refresh schedule will power-off the recovery machines used by the testing/development team, and apply an updated dataset from the protected machines at the set interval schedule. If your testing/development team requires the previously mined dataset to be saved before a refresh, RecoverTrac offers the option to use a Clone Mode instead of a Test Mode. The source host image can be running in a physical server, or as a VM in a hypervisor server such as Citrix XenServer, Microsoft Hyper-V, or VMware ESX Server.

Periodic DR Drill/Rehearsal

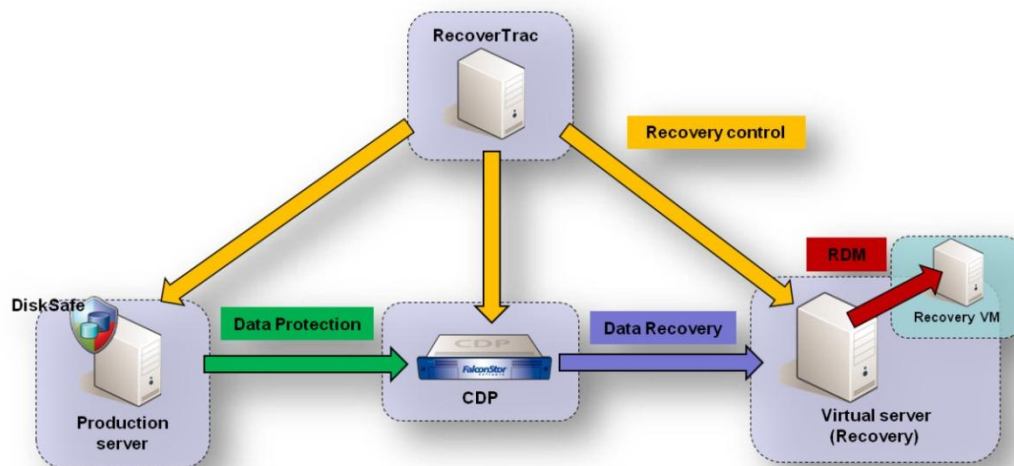
IT organizations should periodically perform DR tests to ensure that processes are correct and environments are working as expected. However, this process is typically very time- and resource-intensive, and is very disruptive to the daily operations of the IT workforce. RecoverTrac addresses these concerns. You can create a recovery job that leverages recovery machines at the DR site that have been allocated for accommodating a site failover in the case of a real disaster, and run the recovery job in Test Mode. The process will be completely automated; it will not impact operations at the production site, nor will it impact data replication between the production and DR sites. Because of its transparency and automation, DR rehearsals can be performed more frequently (even weekly or daily).

Automated Remote Site Failover/Failback for DR (P2P, V2V, or P2V/V2P)

RecoverTrac can orchestrate and automate site failover and failback for all types of applications workloads in the event of a disaster or a planned outage, whether the production host images running in the protected machines are physical servers or VMs.

For each physical protected machine, the paired recovery machine at the DR site (which is defined in the recovery job) can be either a physical machine (same make and model, with nearly identical specs), or a VM running on Microsoft Hyper-V or VMware ESX Server.

For each virtual protected machine, the paired recovery machine at the DR site must also be a VM, but it can run on the same or a different hypervisor. If running on the same hypervisor type, it can be on a different version of the hypervisor.



When performing a site failback, the host images that were loaded onto the paired recovery machines must be recovered on the same initial protected machines. Logically, all data that was modified or updated in the host images while running at the DR site during the site failover period will be replicated back to the production site prior to beginning the failback process. All IP addresses for each host image are properly adjusted to the network environment for each site, and these IP addresses are automatically re-adjusted back to their original value upon failback.

Site Migration for Workload Balancing and Workload Distribution

If your business owns or leases multiple data centers, those data centers may have varying amounts of hypervisor servers, and may have different load-handling capabilities. Seasonally, there could be times when some application workloads (host images) may need to be moved to a different data center location, to accommodate a sudden resource spike. Once that period is over, and because the hypervisor servers of the recovery machine used to host the host image during that period may need to reclaim its resources for other projects, the host image is then migrated off the recovery machine. A new recovery job can be configured to migrate the application workload to either its original hypervisor server at the original site, or to a new destination.

Service Providers Offering DR as a Service (DRaaS)

One of the challenges in selling DRaaS is that the service provider must either convince the customer to use the same hardware as the provider, or adapt to acquire the missing hardware. This hardware can include storage arrays (both sites usually need the same array model in order to perform efficient array-based remote replication), hypervisor platforms (both sites usually need the same hypervisor platform in order to be able to move workload across the sites), and server make and model (for physical workloads in a P2P DR recovery scenario, or to support boot from SAN over Fibre Channel [FC] or iSCSI).


Through RecoverTrac technology, FalconStor CDP addresses those problems, as heterogeneous array replication is handled by built-in delta-based block-level replication, and heterogeneous hypervisor platform conversion is handled via the RTConvert engine of RecoverTrac, which allows a physical or virtual workload to boot on Microsoft Hyper-V or VMware ESX Server, regardless of source machine hardware type (physical or virtual) or hypervisor platform.

VMware Site Recovery Manager Failover/Failback

With RecoverTrac, you can include physical server site failover as part of a VMware Site Recovery Manager recovery plan. RecoverTrac simply creates a VM version (via a recovery machine in scheduled Clone Mode) of the physical server (protected machine) at the production site. This recovery machine, which is a VMware ESX Server VM, is then included in a VMware Site Recovery Manager protection group. In addition, RecoverTrac can operate in tandem with VMware Site Recovery Manager to enable failback.

iSCSI Boot Recovery

If your protected machine is booting the host image via iSCSI, and your recovery job uses a physical server as the recovery machine, RecoverTrac can allow the recovery machine to boot over iSCSI as well. If the recovery machine is a VM, RecoverTrac will adjust the recovery boot disk's geometry to allow it to be attached to the VM's virtual disk adapter, and booted as a standard raw disk.



Conclusion

Many DR solutions focus on replicating data to a remote DR site, leaving IT departments burdened with the complex task of reconstituting servers, applications, network configurations, and replicated data into a functional set of data center services for business continuity. As described in this white paper, RecoverTrac not only replicates data, but stages the recovery of complete services by fully automating the resumption of servers, storage, networks, and applications in a coordinated manner.

RecoverTrac is the first DR automation solution to bring service-oriented recovery to both physical and virtual server infrastructures. RecoverTrac tool automates complex, time-consuming and error-prone failover/failback operations of systems, applications, services and entire data centers, making FalconStor CDP and FalconStor NSS among the most comprehensive disk-based data protection systems available for backup and DR.