the Availability Digest

www.availabilitydigest.com

FalconStor RecoverTrac – Automated Disaster Recovery

June 2012

In our companion article, Recovery-as-a-Service (RaaS),[1] we described emerging cloud services that provide data center backup protection for major outages. The FalconStor RecoverTrac disaster recovery automation tool lets you build your own RaaS cloud.

Downtime and lost data can be catastrophic for a company. In extreme cases, as evidenced by the aftermath of the 9/11 attacks in the United States, companies often cannot survive days of downtime or the loss of their corporate data.

Companies today invest heavily in backup and recovery infrastructure to protect their mission-critical applications. Classic tape backup procedures are unacceptable for truly critical applications since they may take days to recover as the full and incremental backup tapes are loaded. Even worse, days of data – all changes since the last backup – may be lost.

Consequently, companies are moving to hot backup sites in which dedicated servers are booted and ready to be put into service. A synchronized copy of the application and system data is maintained at the backup site via data replication. The backup site can be brought into full production following a production center outage in just a few hours, and only that data caught in the replication pipeline – typically minutes – is lost.

However, failover to the backup site following a production site outage is a complex process and often goes afoul. Many things can go wrong in the heat of the outage crisis as dozens or hundreds of interrelated steps must be accurately executed. Human error, incompatible hardware or software versions, failover process flaws, and network problems are only a few of the challenges that can impede a failover. Because failover is such a complex and risky process, it is seldom tested thoroughly. Rather, companies often depend on faith and hope that they can achieve a successful failover.

The FalconStor RecoverTrac[2] tool solves these problems for hypervisor and physical platforms including Windows and Linux systems. It provides a completely automated recovery solution that can provide recovery times often measured in less than an hour with data loss measured in minutes. Equally important is that failover can be easily tested so that it is assured that it will be successful in a real-life crisis.

[1] HP's Cloud Recovery-as-a-Service (RaaS), *Availability Digest*; June 2012. http://www.availabilitydigest.com/public_articles/0706/raas.pdf

[2] RecoverTrac v2.0: Service-Oriented Disaster Recovery, *FalconStor White Paper*; 2012.

Heterogeneity in a Disaster Recovery Configuration

Data centers have grown up over a long period of time. As a consequence, they are generally populated by a multiplicity of heterogeneous physical servers, virtual servers, storage systems, and operating systems. Today's disaster recovery solutions typically do not support heterogeneity.

For instance, VMware's Site Recovery Manager supports production systems that are running virtual machines on VMware ESX servers only. Physical servers are not supported, nor are other hypervisors. The same can be said for Citrix' Site Recovery for XenServer.

FalconStor's RecoverTrac solves the heterogeneity problem to a great extent. It supports virtual-to-virtual (V2V), physical-to-physical (P2P), and physical-to-virtual (P2V) failover configurations. For V2V failover, the production system can be running virtual machines (VMs) on any certified hypervisor (VMware, Microsoft Hyper-V, or Citrix XenServer). RecoverTrac is capable of failing these VMs over to either VMware or Hyper-V virtual environments.

For P2V recovery, RecoverTrac recovers the physical server to an appropriately configured virtual machine in the recovery center. P2P failover requires that there be a physical server of the same make and model in the recovery center.

RecoverTrac's Replication Channel

In order to bring up a physical backup server or a VM in the recovery center, a current copy of the application's database and its operating system image must be available in the recovery center. This is accomplished via the FalconStor Continuous Data Protector (CDP) replication engine. Replication is heterogeneous – FalconStor CDP can replicate from any supported database or application. Supported databases include Oracle, DB2, Informix, Ingress, SQL Server, MySQL, and Sybase. Supported applications include SAP, Lotus Notes, and Exchange, among others. The replication target at the recovery site can be any supported fibre channel or iSCSI storage array.

The FalconStor CDP replication engine (Figure 1) comprises a pair of Windows-based appliances, a FalconStor CDP appliance at the production center and a FalconStor Network Storage Server (NSS) at each of the recovery centers to which the data is to be replicated.

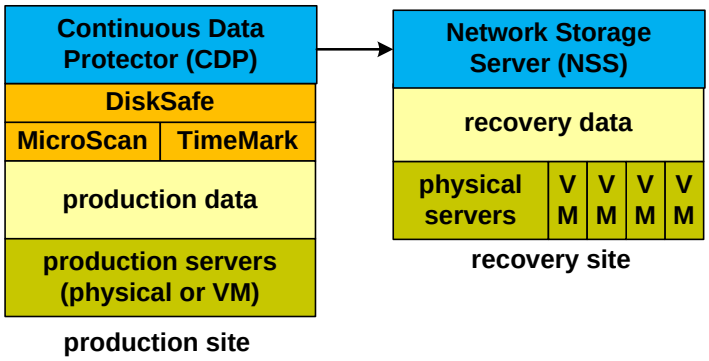

FalconStor CDP Replication Engine
Figure 1

FalconStor CDP replication can either be continuous or scheduled. In addition, point-in-time snapshots can be taken and replicated to the recovery site along with the current data changes. FalconStor CDP ensures that the current state of the databases and the snapshots at the recovery site are always application- and transaction-consistent.

FalconStor CDP supports one-to-one, one-to-many, many-to-one, and many-to-many replication streams. For instance, one recovery center can act as the backup for multiple production sites. Alternatively, the backup of one data center may be distributed among several recovery centers for load-balancing purposes.

RecoverTrac can operate by itself or in conjunction with other disaster recovery configurations. For instance, it can provide P2P failover for VMware Site Recovery Manager.

The FalconStor CDP appliance at the production center is fed by an agent resident on each protected physical or virtual production machine and supports both direct attached storage and storage arrays. The agent sits between the applications and the file system of the host server. The agent comprises:

- *MicroScan*, which monitors for disk writes that will be replicated. It deduplicates the write by selecting only the sector or sectors out of the disk block that have changed. Typically, only 512 bytes out of a 32K-byte disk block are replicated.
- *TimeMark*, which creates point-in-time application-consistent snapshots of the data for replication.
- *DiskSafe*, which transfers the data to be replicated to the FalconStor CDP appliance at the production site.

The FalconStor NSS appliance at the recovery center writes the replicated data to the particular SAN storage arrays in use at the recovery center.

RecoverTrac Failover and Fallback

Given a synchronized copy of the production databases at the recovery center, RecoverTrac provides automatic or manual failover of a failed production environment to the recovery data center. Once the production environment is returned to service, RecoverTrac seamlessly manages the fallback process to restore normal operations at the production center.

RecoverTrac runs in its own server, as shown in Figure 2. It has communication and control links with all of the protected servers at the production site, the backup servers at the recovery site, and the FalconStor CDP replication appliances.

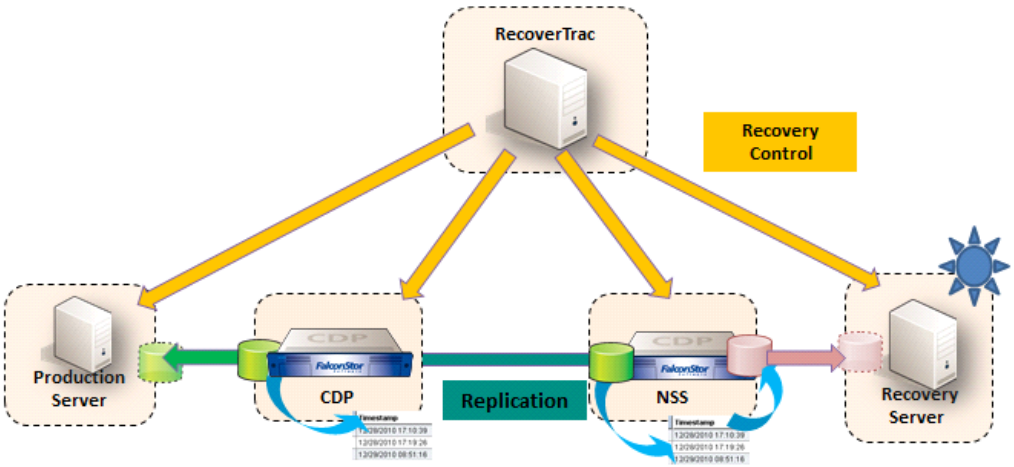

RecoverTrac Outage Management
Figure 2

Failover

RecoverTrac provides fully automated recovery or manual recovery through a management interface. The recovery decision is often a manual process as it must be determined whether it would be faster to try to recover the failed production systems or to failover to their backup systems.

RecoverTrac automates the entire recovery process. When RecoverTrac is notified of a production outage requiring failover to the recovery site, it will boot the required VMs and physical servers with their operating system images at the recovery site, and it will mount the pertinent application databases. It

changes the IP addresses used by the users so that they are now routed to their recovery backup systems, and production processing proceeds.

Failed virtual machines can be failed over to either a Microsoft Hyper-V or a VMware ESX hypervisor. Likewise, failed physical machines can be failed over to VMs running under either of these hypervisors. Alternatively, physical machines may be failed over to other physical machines of the same make and model.

RecoverTrac can fail over to the current replication state at the time of the outage, or it can recover to one of its snapshots at a specified point in time. This latter capability is useful, for instance, if a database should become corrupted.

Fallback

When the production environment is returned to service, applications are returned to the production site via the fallback process. RecoverTrac fully automates the fallback process.

Fallback is essentially the reverse of failover. The first step is to use the FalconStor CDP replication engine to bring the production databases and system images into synchronization with the current recovery environment. The production physical servers and virtual machines are rebooted from their system images, and application databases are mounted. Finally, the user IP addresses are redirected to the production environment. Normal processing at the production site is now restored.

Via RecoverTrac Re-Home capability, any changes made to the processing environment while in recovery mode can be made to the production systems prior to the completion of the fallback process.

Other Uses

RecoverTrac can be used for a variety of other purposes:

- *Testing DR failover*: Perhaps one of the most important uses for RecoverTrac beyond disaster recovery is the ability to test failover in a test environment without impacting production. In this case, the recovery environment is brought into operation just as it would be in a real emergency, but user IP addresses are not changed. Users continue to use the production environment. The recovery environment can be tested, and then the recovery databases resynchronized to return it to backup mode.

- *Local bare metal recovery*: A recovery machine located at the production site can run a recovery job, automatically moving an application workload from a production machine to the recovery machine.

- *Remote Test/Development Team Data Refresh*: Special recovery jobs can be scheduled to synchronize test and development machines at remote sites for data mining, testing, and analysis. The remote machines can be synchronized with the current data or with a specified snapshot.

- *Site migration for workload balancing and workload distribution*: Recovery machines can be activated at remote data centers to move some workload from the current production servers to the recovery servers during periods of peak processing activity. When the workload subsides, the workload can be returned to the production center.

- *Service providers offering Recovery-as-a-Service (RaaS)*: RecoverTrac solves the heterogeneity problem facing RaaS service providers. With RecoverTrac, heterogeneous production systems, heterogeneous storage arrays, and heterogeneous hypervisors are supported.

Summary

RecoverTrac supports fully automatic or manual V2V, P2V, and P2P failover and fallback in heterogeneous environments. Backup services can be provided for any supported application and database. The production and recovery hypervisors and storage arrays do not need to be the same.

Recovery can be to the current operational state or to a prior snapshot. All replication states are guaranteed to be application-consistent and transaction-consistent.

RecoverTrac automatically redirects user links to the recovery environment as part of failover and back to the production environment during the fallback process. However, failover can be easily tested with no risk to the production environment in RecoverTrac test mode as users are not redirected away from the production environment.

RecoverTrac logs all recovery operations for a complete audit trail.