

AUTORECOVER:

Automated Service-Oriented Disaster Recovery

Contents

Abstract.....	3
Introduction	3
Standard Disaster Recovery Solutions	3
Addressing Common Challenges.....	4
OVERVIEW AND BENEFITS	6
Any Service.....	6
The need for AutoRecover.....	6
Service-Oriented Data Protection	6
Any Time.....	7
Rapid Recovery.....	7
Automated Failover and Failback	7
Any Place.....	7
Any-to-Any Recovery.....	7
How Does AutoRecover Technology Work?.....	8
Production Site.....	8
Disaster Recovery Site	9
Key Features	10
Local Bare Metal Recovery (P2P, V2V, P2V	12
Remote Test/Development Team Data Refresh	13
Periodic Disaster Recovery Drill/Rehearsal	13
Automated Remote Site Failover/Failback for Disaster Recovery (P2P,V2V,P2V,V2P).....	14
Site Migration for Workload Balancing and Workload Distribution.....	15
Service Provider Offering Disaster Recovery as a Service	15
VMware Site Recovery Manager Failover/Failback.....	15
iSCSI Boot Recovery	15
Conclusion	16

Abstract

Protecting your data means nothing if you cannot recover it. One of the largest data center challenges is recovering data smoothly and quickly after your system goes down. Downtime can be caused by data corruption, power outage, or natural disaster. Resuming operations at a disaster recovery (DR) site, whether the outage was a planned event, such as a scheduled site migration, or an unplanned event, such as an accident, requires careful preparation.

While planning your DR strategy may take months, execution of your DR plan must occur in minutes.

Do you want to rely on the ability of your data center staff to arrive at the designated DR location to manually recover your production environment? Hoping your staff can respond in an emergency and perform the numerous coordinated steps to restore your hardware, network, storage, and applications is dangerous and risky.

Take the danger and risk out of disaster recovery with X-IO's iglu blaze AutoRecover.

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, corruption, human error, equipment failure, or a complete site outage due to power loss or as a result of 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 can take months, but the execution of the plan needs to occur within minutes. During these precious minutes, all of the teams involved are 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, or a number of other factors could delay the system or site recovery. For those without the right tools, this herculean effort can be both risky and unpredictable.

X-IO believes that disaster recovery should not be this complex, and has a great deal of expertise and experience in this sector. For this reason, X-IO provides AutoRecover™, an automated disaster recovery tool as a key feature of iglu blaze. The AutoRecover tool automates complex disaster recovery tasks, bringing service-oriented recovery to both physical and virtual server infrastructures. This technical white paper discusses the concerns associated with disaster recovery, and explains in detail how AutoRecover technology addresses a full range of recovery challenges.

Standard Disaster Recovery Solutions

Various solutions from various vendors have been designed to orchestrate the disaster recovery workflow. VMware, for example, has designed its Site Recovery Manager (SRM) product to enable simple, one button-recovery and site failover execution. VMware SRM leverages array-based replication to send

protected data to the disaster recovery 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 SRM, 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. Because the solution operates at the site or array level, the granularity of individual server failover and recovery is compromised. While site failover is automated, site failback (the ability to move a failed-over workload from the disaster recovery site back to the production data center) is a complex and often manually intensive process that reintroduces significant risk and uncertainty.

Microsoft has taken a slightly different approach to disaster recovery 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 Microsoft 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 provides VM-based applications that are not MSCS-aware with high availability (HA) 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 can an organization orchestrate a successful recovery? And what if the disaster recovery site and the production data center operate on different computing and storage platforms?

Addressing Common Challenges

When it comes to data recovery, and in particular, disaster recovery planning, organizations face a common set of concerns. Do any of these sound familiar?

- Our budget has been cut, but we need a good disaster recovery plan. What can we do?
- How do I protect physical machines? Do I need to recover to an identical physical machine at the DR site? Can I use a VM as the disaster recovery standby machine?
- I have a handful of business-critical applications that need to recover in minutes, not hours. Is there a rapid recovery solution available for select systems and applications?
- I cannot implement VMware SRM because of various hardware/environmental reasons. We need complete recovery coverage.

- Most solutions I have looked at only guarantee crash-consistent data. Can I make sure that after failover my tier-one database will be mounted cleanly without wasting time repairing inconsistencies?
- During a site outage, I may want to temporarily move my protected workloads, physical machines, and VMs to the disaster recovery site. Can this process be automated? How can we automate the entire failback process back to the production site?

Included with the iglu blaze solutions, AutoRecover technology automates the entire disaster recovery process in any type of data center or heterogeneous environment. The AutoRecover tool addresses a full range of recovery needs, from local data recovery such as bare metal recovery, to remote data recovery, with full site failover and failback. This includes like and unlike physical-to-physical (P2P), virtual-to-virtual (V2V), physical-to-virtual (P2V), and virtual-to-physical (V2P) recovery for any certified hypervisor or physical platform, even with multiple network segments.

Organizations may have additional concerns regarding automated disaster recovery and how AutoRecover technology will respond in various situations:

- If one of my production servers is a physical system and I want to move it offsite as a VM, can I move it back and convert it to run again on the original physical machine after failback?
- Are VMware ESX(i) Server 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 possible?
- I want to rehearse/test a 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 we seldom perform these rehearsals. Does the AutoRecover tool allow us to automate disaster recovery testing?
- From the production data center, our engineering team wants our Microsoft SQL servers to be replicated and tested in one disaster recovery site, while the accounting department wants the file/print servers to be replicated and restored in another site. Can AutoRecover technology handle one-to-many, many-to-one, and many-to-many site configurations?
- We need a solution that is 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 configuration, including the recovery jobs, to a different server?
- If a real disaster occurs, I need to be up and running very quickly. Does AutoRecover technology allow us to quickly recover a machine, and can machines be recovered in parallel?

By enabling all of the above mentioned scenarios, AutoRecover technology makes disaster recovery planning, testing, and execution as simple as possible. AutoRecover 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 as a VM.

The next section discusses how AutoRecover technology operates in such situations.

Overview and Benefits



The AutoRecover tool delivers automated service-oriented disaster recovery for all your data. AutoRecover is available for 'Any Service, Any Time, Any Place.' We will explore each of these elements.

Any Service...

The need for AutoRecover

Most organizations today rely on multiple applications, server platforms, networking protocols, storage systems, and other IT resources to operate their business. Unfortunately, service interruptions can come from sources such as hurricanes and other natural disasters, local and regional power grid failures, and human errors or malicious actions. Rebuilding this infrastructure at a remote data center is challenging, and with costs and losses that average well over \$100,000 per hour. Many vendors and customers mistakenly believe that just having a secondary instance of data constitutes disaster recovery. However, backups, snapshots and/or replication alone do not fulfill all of the complex requirements needed to deliver on the promise of disaster recovery.

Service-Oriented Data Protection

Service-oriented data protection matches data protection to the way data centers actually manage IT. Administrators are required to deliver a specific service, such as email, web portal, or sales force automation. These can be very complex systems, such as HTTP servers talking to collaboration software or databases while using Linux and Microsoft Windows. Customers need to apply backup policies such as disaster recovery, failover, failback, archive, and retention to the entire service and not only at the file and

block level. Although they still do file and block backup, they can now protect all of the elements of a specific service as one interrelated group. The AutoRecover tool is a key part of this service-oriented data protection offering from X-IO.

Any Time...

Rapid Recovery

In a real emergency, the only thing that matters is getting a down data center back up and running. The AutoRecover tool empowers organizations to recover data in minutes, not hours or days. Not only can it recover a single machine in minutes; it can perform up to five parallel recovery jobs by default. This number can be increased when using faster hardware.

Automated Failover and Failback

The AutoRecover tool can start up target hosts (physical or virtual) and applications at the recovery site, shut down affected hosts at the primary site, and reverse the direction of remote replication, maintaining data protection when the primary data center has storage functionality. If the primary site failure is server-specific, this dynamic reversal of direction is essential for failback to the primary site. Since most failures are man-made and short-term in duration, failback is vital for comprehensive, automated disaster recovery.

Any Place...

Any-to-Any Recovery

AutoRecover technology provides the administrative flexibility that is desperately needed in case of an actual disaster. Administrators can use any combination of supported physical and virtual machines for any-to-any recovery: P2P, P2V, V2V, or V2P, including similar or dissimilar machines. The AutoRecover tool features seamless integration with VMware and Microsoft Hyper-V. Intuitive wizards help create hypervisor servers and virtual machines (VM) at any location.

Not only can the AutoRecover tool become the sole orchestrator for all of an organization's failover/failback needs; it can integrate with existing environments that already have a disaster recovery automation solution in place, such as VMware SRM. As noted, VMware SRM was conceived from the beginning 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 disaster recovery protection. The remaining 15-20% consists of physical servers that cannot be virtualized for one reason or another.

Think about those robust, highly-demanding x86 Microsoft Windows or Linux servers that the application administrative team simply refuses to virtualize. The AutoRecover tool can complement VMware SRM and take care of these remaining and highly important servers. In fact, for the x86 physical servers, AutoRecover technology can offer multiple recovery choices for disaster recovery site failover:

- If VMware SRM has been installed, the AutoRecover tool can stage a P2V conversion of the physical server at the production site to create a standby VM instance that the VMware solution can fail over as part of its recovery plan. This extends VMware SRM capabilities beyond protecting and failing over

VM workloads alone. If VMware SRM is not installed or available, the AutoRecover tool can be used to create recovery plans of VMware environments. AutoRecover technology can replicate changed data on each VM. When recovery is needed, it can perform both failover and failback of VMs in a single job.

- AutoRecover technology can also leverage VMware APIs for on-demand VM creation. There is no need to create target VMs that remain unused. This reduces VMware licensing costs and overhead at the disaster recovery site.
- The AutoRecover tool can help automate the site failback process. AutoRecover technology can stage the entire site failover, and keep the physical workload at the disaster recovery site.
- The AutoRecover tool can stage the entire site failover, including performing a P2V conversion at the destination disaster recovery site, to turn the physical workload into a VM. The VM can be hosted on VMware or Microsoft Hyper-V.

How Does AutoRecover Technology Work?

The AutoRecover tool works in conjunction with iglu blaze. The data protection solution is included in iglu blaze that provides unified backup and disaster recovery, providing fast recovery to any known good point in time. While iglu blaze protects application servers and replicates data across data centers, the AutoRecover tool provides additional value to the recovery process through disaster recovery automation.

The AutoRecover tool can perform both local recovery (bare metal recovery) and remote recovery of servers. For remote recovery, AutoRecover can handle many-to-many site mappings. For example, two data centers, one in Boston and one in Chicago, could replicate data 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 disaster recovery sites in different cities based on load capacity.

For the following example, we will assume two sites: a production site and a disaster recovery 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, an iglu blaze CDP protection package is installed. For Microsoft Windows systems, this package includes CDPagent™ technology from X-IO, as well as application-specific snapshot agents. A complete list of supported applications and databases can be found at <http://www.X-IO.com/certification-matrix/applications><http://www.X-IO.com/certification-matrix/applications-and-databases>.

An iglu blaze is installed locally with enough storage capacity to track changes for snapshots and replication purposes. This local array backs up data from the protection application workload using changed-block tracking technology called RepIOptimize™, with a flexible scheduling option that offers

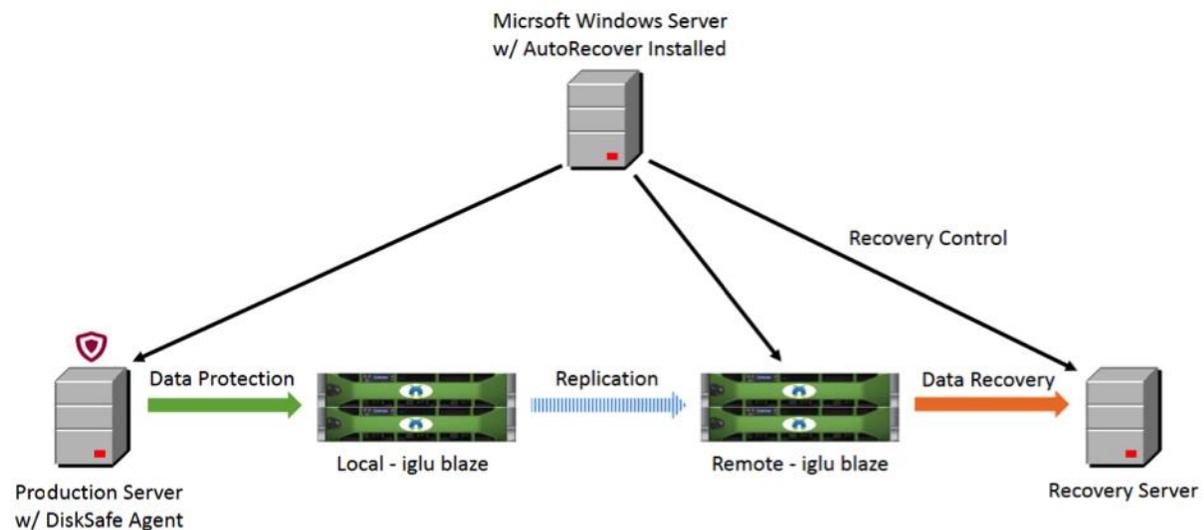
continuous or periodic protection modes. This means that a backup schedule can be as frequent as every 10 or 15 minutes, instead of restricted to once per day). The resulting backup points are called TimePoint® snapshots, which are point-in-time quiesced with full application consistency. The snapshots for each application workload can be replicated remotely to iglu blaze at one or several disaster recovery sites.

For local recovery, install the AutoRecover tool on a Microsoft Windows server at the local site. This orchestrates the recovery of workloads by communicating with the iglu blaze and automating the local data restore process.

Disaster Recovery Site

At the recovery site, physical servers are installed with the same, similar, or even dissimilar hardware specifications as the physical server in the production site. Alternatively, a supported hypervisor server capable of running VMs with acceptable performance may be present.

Another iglu blaze is deployed as a replication target. 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.



AutoRecover technology is installed on a Microsoft Windows Server at the disaster recovery site for use should the production site become unavailable during site outages. This AutoRecover tool will orchestrate disaster recovery for all of the protected workloads coming from the production site(s).

The two iglu blaze storage arrays, local and target, are basically the backup units. The iglu blaze CDP software protection package, installed in each protected workload, performs the scheduling for point-

intime, transactionally consistent snapshots, as well as delta data movement, with changed-block tracking enabled from the primary disk to the local appliance.

Iglu blaze' management interface allows users to perform the initial configuration for replication schedule and site pairing. Beyond the initial deployment, IT teams can use this console if they wish to perform manual recovery of data by manually creating a mountable snapshot, assigning that synthetic restored disk as a drive letter or mount point to a designated server, and copying the data back. However, it is not necessary to perform manual recovery, since the AutoRecover tool automates every step, while providing many other compelling features.

Once the two iglu blaze storage arrays have been configured to replicate between one another, and the application workload has been protected with the CDPagent agent and optional snapshot agents, organizations can then configure the AutoRecover tool to authenticate to both iglu blaze arrays, so that either one can perform recovery tasks and jobs for both sites.

The concept behind AutoRecover technology is simple: AutoRecover technology carefully de-couples the machines, which define the hardware (virtual or physical), from the application workloads or host images, which define the software (operating system and application data). The AutoRecover tool 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

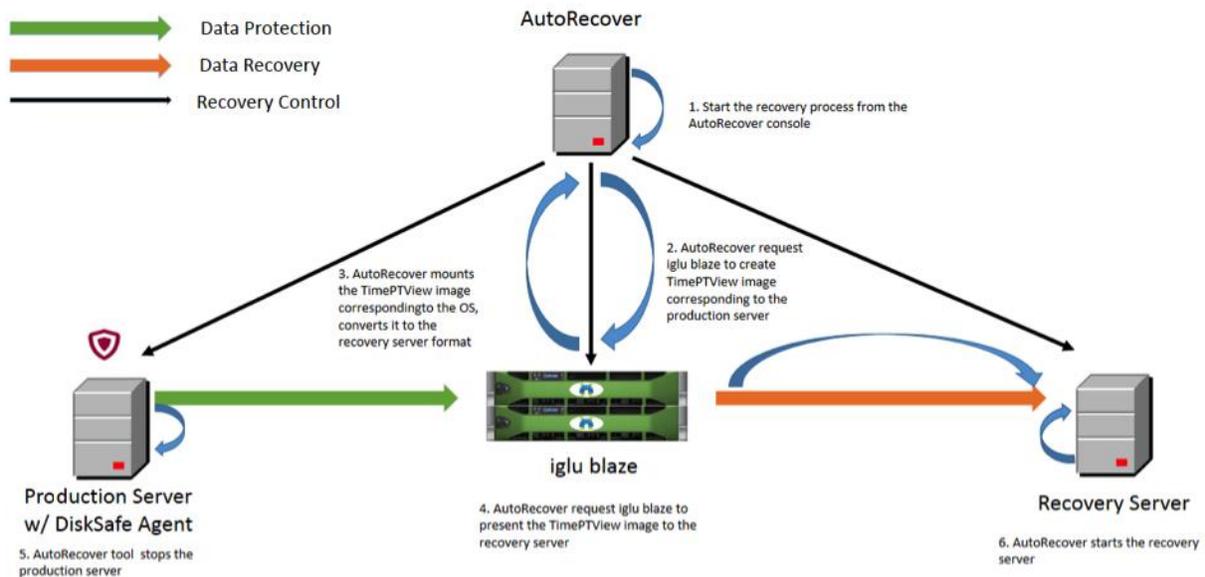
- Recover data in minutes, not hours or days. This is important because, in a real data emergency, the only thing that matters is getting the data center up and running. The AutoRecover tool can not only recover a single machine in minutes, but can perform up to five parallel recovery jobs (or more with a more powerful hardware configuration).
- iglu blaze leverages ReplOptimize technology for changed-block tracking, which allows for extremely low-level, disk sector (512-bytes) delta data replication, as well as compression and encryption, for the most efficient and secure sub-block-based replication available. This considerably reduces the recovery point objectives (RPO), as replication can be performed more frequently and with less network bandwidth utilization. Application workloads (host images) are recovered with full transactional consistency due to snapshot agents supporting database applications. This allows the AutoRecover tool to perform recovery with the smallest possible recovery time objectives (RTO), and with zero data loss.
- 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 like or unlike physical and virtual machines.
- P2P site failover and site failback support, from supported physical server models to like or unlike physical server models.

- V2V site failover and site failback support, in any combination of hypervisor platforms. Regardless of how or where an error occurs, automated failback is vital for continuous data and service availability. AutoRecover technology automates failover and failback, shutting down affected hosts at the primary site, booting up target hosts and applications at the recovery site, and reversing the direction of replication.
- Intelligent Failback replicates all data back to the production site, including both the recovered data and any new data created at the recovery site.
- 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 an iglu blaze resource, the AutoRecover tool 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 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 X-IO NSS and related snapshot agents from X-IO, the new recovery machine must get its snapshot notifications from the iglu used for disaster recovery. Finally, if the recovery machine is connected to iglu data disks over iSCSI, it also must be reconfigured to reconnect to the iSCSI target at the target appliance located at the disaster recovery site (not the iSCSI target at the production site). The AutoRecover tool automates this process to ensure a smooth migration/transition to the recovery site, and to prepare for a smooth site failback.
- In addition to “Recovery Mode,” which is used in an actual disaster recovery scenario to restore machines, AutoRecover technology offers two additional recovery methods:
 - **Clone Mode**, to branch off to a machine independent of its source. Useful for short- or long-term testing/development projects, or for cloud service deployments from templates.
 - **Test Mode**, to rehearse disaster recovery plans. You can schedule disaster recovery rehearsals and pre-stage environments to speed up recovery during actual disaster events. You can test and validate disaster recovery plans and environments without impacting production machines, applications, or data. Network fencing enables separation of production and non-production environments to test disaster recovery.
- Ability to change IP address of each host image during a recovery job execution, to accommodate the network scheme of the disaster recovery site. This is needed if a disaster recovery 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 the AutoRecover tool retains all of the machine information for all of the sites in its database.
- Ability to execute recovery jobs with any available TimePoint snapshot. Users 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 disaster recovery orchestration (to accelerate the site failover process).
- Deep VMware integration including VMware vCenter Server 5.x support and VMware vCenter Cluster Aware support. For example, the AutoRecover tool can recover 50 hosts to a VMware cluster, and VMware will automatically distribute the load across the cluster. This feature leverages VMware APIs for on-demand VM creation to reduce VMware licensing costs and overhead at the disaster recovery site. It provides simplified management with no need to create target VMs and have them remaining unused. In addition, it includes optimized VMware calls and rescan logic. Recovery jobs can now automatically set/update the VM Network settings, supporting network fencing to separate production and test networks. Furthermore, AutoRecover technology is VMware Cluster aware to support load balancing and HA configurations
- An efficient recovery engine allows for more automation logic and complex recovery grouping and segmentation (job chaining).
- Capable of chaining several jobs together, allowing jobs to have any level of granularity required by user. Enables easy cleanup of replication when in failover.
- Authoritative restore for Active Directory. Eliminate all manual steps by streamlining VM and OS commands during recovery to save time and reduce errors. Supports authoritative restore of Domain Controller for isolated testing or large-scale disasters.
- Streamlined console integration. The AutoRecover console manages all aspects of recovery. Automated failback allows customers to resume normal operations back at the original site. ReplOptimize technology, Convert, and driver verifications ensure proper resumption of machines and services/applications.
- Command Line Interface (CLI) for additional recovery flexibility. This includes job creation, change, training, scheduling, grouping, pre/post scripts, and other functions.
- Highlights delta changes at failover or failback site to alert of potential issues before the conversion happens. Delta alerts allow IT managers to quickly spot potential issues and fix them before conversion/recovery (network configurations, port mapping, driver issues, and so on).
- An AutoRecover machine profile utility automates the collection of drivers/configurations to simplify creating machine profiles. Enter and validate environment information to ensure accuracy the first time. The AutoRecover tool centrally maintains relationships and groups simplifying job creation, scheduling, and modifications, enabling drag-and-drop simplicity.

Local Bare Metal Recovery (P2P, V2V, P2V)

The AutoRecover tool allows organizations 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 can be a VM or physical machine. If the protected machine is a supported physical server, then the user can define a recovery job and select either a VM or a supported physical server as the recovery machine.



Remote Test/Development Team Data Refresh

Testing and development teams at various sites across an organization often need to access copies of production data to perform data mining, testing, and analysis. The AutoRecover tool allows IT teams to schedule special recovery jobs, which can bring recovery machines online at remote testing/development centers. These recovery machines contain a carbon copy of the protected machine from a specific point in time. This TimePoint snapshot can be selected from all of the snapshot states that were replicated, or if needed, a user can specify a schedule to perform a dataset refresh. 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 the testing/development team requires the previously mined dataset to be saved before a refresh, the AutoRecover tool offers the option to use a Clone Mode instead of a Test Mode. The source host image runs in a physical server, or as a VM in a hypervisor server such as Microsoft Hyper-V or VMware ESX(i) Server.

Periodic Disaster Recovery Drill/Rehearsal

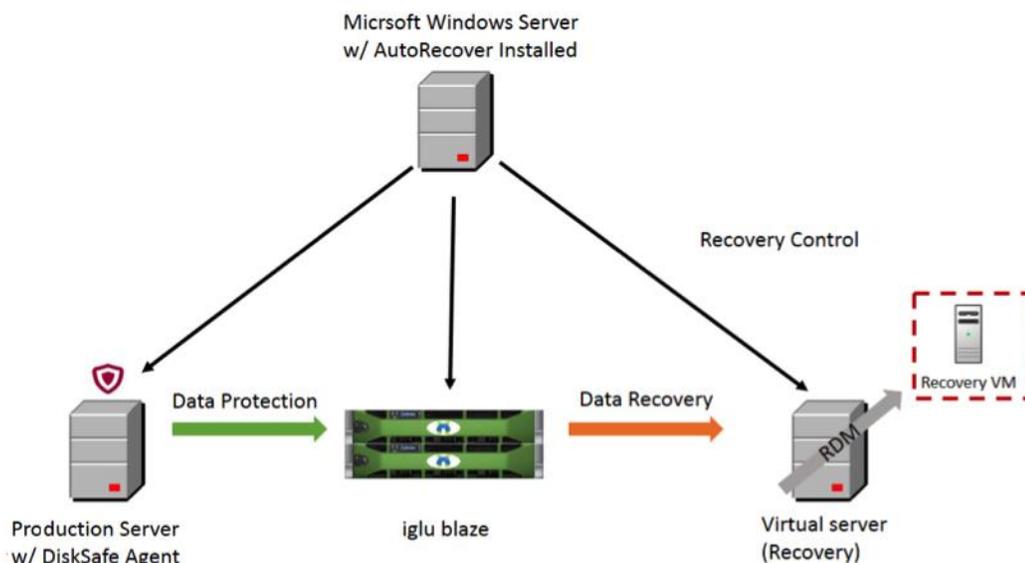
IT organizations should periodically perform disaster recovery tests to ensure that processes are correct, and environments are working as expected. However, this process is typically very time- and resource-intensive, and disruptive to the daily operations of an IT workforce. AutoRecover technology addresses these concerns. A user can create a recovery job that leverages machines at the disaster recovery 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 is completely automated, and does not impact operations at the production site or replication between sites. Because of the transparency and automation of AutoRecover technology, disaster recovery rehearsals can be performed more frequently – as often as weekly or daily.

Automated Remote Site Failover/Failback for Disaster Recovery (P2P,V2V,P2V,V2P)

The AutoRecover tool can orchestrate and automate site failover and failback for all types of application 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 disaster recovery site (which is defined in the recovery job) can be either a physical machine (similar or dissimilar specs), or a VM running on Microsoft Hyper-V or VMware ESX(i) Server.

For each virtual protected machine, the paired recovery machine at the disaster recovery site can be a physical machine or VM (running on the same or 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 disaster recovery site during the site failover period will be replicated back to the production site prior to beginning the failback process. By using Intelligent Failback you can plan resynchronizations back to the production environment in advance of the failback process, thereby minimizing the time required to fail workloads back to the production environment. 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 Provider Offering Disaster Recovery as a Service

One of the challenges in selling disaster recovery as a service 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 disaster recovery scenario, or to support boot from SAN over FC or iSCSI).

Through AutoRecover technology, iglu blaze 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 process of the AutoRecover tool, which allows a physical or virtual workload to boot on Microsoft Hyper-V or VMware ESX(i) Server, regardless of source machine hardware type (physical or virtual) or hypervisor platform.

VMware Site Recovery Manager Failover/Failback

With the AutoRecover tool, a customer can include physical server site failover as part of a VMware SRM recovery plan. AutoRecover technology 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(i) Server VM, is then included in a VMware SRM protection group. In addition, the AutoRecover tool can operate in tandem with VMware SRM to enable failback.

iSCSI Boot Recovery

If a protected machine is booting the host image via iSCSI, and a recovery job uses a physical server as the recovery machine, AutoRecover technology can allow the recovery machine to boot over iSCSI as well. If the recovery machine is a VM, the AutoRecover tool 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.

In testing with Login VSI, storage performance (IOPS) required during the user login phase was two times the "steady-state" workload. Considerable thought should be given to user behavior patterns, as to how many users will be concurrently initiating connections to the desktop environment. Login VSI is capable of adjusting the rate at which users log into the environment for the test, and is an excellent method for

exploring the performance of the solution during various login activity levels. In the testing X-IO performed, user login rates of 16/min, 25/min, and 33/min were tested. The single greatest impact to the login times was the high compute (CPU and RAM) utilization of the UCS blade servers when the majority of the users had connected to the solution.

Conclusion

Many disaster recovery solutions focus on replicating data to a remote 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, the AutoRecover tool from X-IO 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.

The X-IO AutoRecover disaster recovery automation solution brings service-oriented recovery to both physical and virtual server infrastructures. The AutoRecover tool automates complex, time consuming, and error-prone failover/failback operations of systems, applications, services, and entire data centers, making it a mainstay among the most comprehensive disk-based data protection systems available.

Interested in learning more or starting a proof of concept?

Contact us today at fastforever@x-io.com,

+1 719.388.5500, +1 800.734.4716 or +44 20 3239 4669

9950 Federal Drive, Suite 100 | Colorado Springs, CO 80921 | U.S. >> 1.866.472.6764 | International. >> +1.925.298.6061 ;
www.xiostorage.com

X-IO, X-IO Technologies, ISE, iglu blaze, IAF and CADP are trademarks of Xiotech Corporation. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. © Xiotech Corporation. All rights reserved.

Document Number WP-0015-20160414