



How to Protect and Recover VMware Environments

Introduction

The use of virtual server technology is rapidly becoming a mainstream computing practice. Server consolidation projects that deploy multiple “virtual” servers on a single physical server take advantage of the fact that many physical servers today are over-resourced to accommodate growth. By running virtual servers at much higher utilization rates, enterprises are able to cut over-provisioning waste and make much more efficient use of their computing resources. In many cases this leads to significant cost savings across many areas, including ease of management, lowered power and cooling costs, decreased physical infrastructure requirements, and faster, easier deployment of new systems. Gartner Inc. forecasts that by 2012, 48% of all x86-based workloads will be running on virtual servers.

As the provider of the first commercially available virtual server platforms, VMware owns the lion’s share of the virtual server market. A survey completed by IDC in late 2009 indicated that 67% of all survey participants were using virtual server technology from VMware. Relative to competitive offerings, VMware has a more mature solutions portfolio in the virtual server market. Because of their relative maturity, when it comes to deploying production applications on virtual servers, enterprises choose to deploy on VMware more often than on any other platform. Because of VMware’s widespread use for production application environments, recovery is increasingly becoming a critical consideration for VMware customers. As the performance capabilities of VMware virtual servers continues to increase, more and more mission-critical applications will be deployed in these environments, making recovery even more important.

As VMware customers consider how to deploy recovery solutions for their VMware environments, there are other industry trends which need to be taken into account. Data is growing at unprecedented rates, and IDC has forecast growth rates of 50% - 60% to continue for at least the next 4-5 years. Evolving business and regulatory mandates are driving increasingly stringent recovery requirements, both in terms of minimizing data loss on recovery as well as in shortening recovery times. And environments are becoming increasingly complex. Consider that most shops today have a mix of physical and virtual servers, running one or more Windows, Linux, or Unix distributions, have a variety of different storage architectures deployed (SAN, DAS, and/or NAS), and are running applications from a variety of different vendors. These trends, combined with the increasing penetration of virtual server technology, are posing significant challenges for incumbent recovery products as administrators try to resolve backup window, recovery point objective (RPO), recovery time objective (RTO), and recovery reliability issues across their heterogeneous environments.

Defining Recovery in VMware Environments

Most enterprises' businesses are driven by a set of critical processes that rely heavily on information technology (IT) infrastructure. In order to ensure that these business processes can continue to run, enterprises must have recovery plans in place that allow them to rapidly and reliably restore the data and application services that drive them, regardless of the reason for the outage. Just being able to recover data is not sufficient, since data is useless without the applications that make sense of that data. For that reason, most enterprises have plans in place to recover both data and applications. But the legacy processes on which businesses have relied to do that are now imposing costs on them which are becoming increasingly difficult to bear. Common "recovery" problems VMware customers can experience when deploying virtual servers can include:

- Data protection operations (i.e. backup) impose too much overhead on production environments
- Too much data is being lost when recoveries are required because "backups" cannot be performed frequently enough
- It takes too long to recover data and/or application services when a recovery is required
- There is too much risk in recovery operations, due either to the poor reliability of tape-based infrastructures or because recovery operations rely too much on manual rather than automated processes
- Too many different products need to be licensed and integrated to provide a complete solution for VMware recovery, leading to higher costs and increased complexity

Most enterprises already have a backup solution in use for physical servers, and it is common when first deploying VMware to just put backup agents on each virtual server. Since much of the savings in a server consolidation project comes from the fact that virtual servers are provisioned at much higher ratios than physical servers, there is generally not enough headroom available on virtual servers to run backup agents without unduly impacting production application performance. This is in fact such a widespread problem that VMware introduced a special backup API in VMware Virtual Infrastructure 3, called VMware Consolidated Backup (VCB), to help address it. VMware provided this hypervisor-level API (within ESX Server) to allow disk-based snapshots to be taken for backup purposes without imposing much impact on virtual server performance. When VMware introduced vSphere 4, the name of this API was changed to the vStorage API.

Using a hypervisor-level API for backup introduced two new issues however. First, customers could not always reliably create an application-consistent recovery point using it since it did not allow direct, full access to application snapshot APIs like

The Importance of Application-Consistent Recovery

Most administrators want to recover from application-consistent recovery points because this supports the fastest, most reliable recovery option for any given application environment. Crash-consistent recovery points can generally produce reliable recovery, but recovering an application from them can take significantly longer.

When applications had to be shut down to be backed up, making a copy of an application-consistent recovery point was not an issue because shutting an application down put it into an application-consistent state. But as backup windows shrank due to the widespread use of 24x7 operations, backups had to be performed on-line when applications were in use, and ensuring application-consistent recovery points became an issue. Application vendors of data services products like databases and file systems generally provided an API that could be used by third party applications to enable the creation of application-consistent recovery points even while an application was in use. Popular application snapshot APIs in use today include Windows VSS, Oracle RMAN, and others. To gain the full access to application snapshot API controls needed to reliably create application-consistent recovery points, a third party agent of some kind (backup agent, snapshot agent, etc.) generally had to interact directly with the application being "backed up."

Windows Volume Shadowcopy Services (VSS), Oracle RMAN, and others. The API did allow them, however, to create disk-based snapshots that could be used for recovery (or off host backup) purposes. And second, using the hypervisor-level API required custom scripting that created a VMware-specific solution. For customers that also had physical servers, and possibly even other virtual server platforms, other data protection solutions would be required in addition.

In addition to local data recovery, there were two other recovery issues to address: recovering virtual servers and the applications on them, and disaster recovery (DR). VMware products like vMotion, vHA and vFT could be used to address virtual server "high availability" (HA) requirements, allowing virtual servers to be migrated from one physical server to another as needed as long as those physical servers were using the same networked storage. VMware tools like Site Recovery Manager (SRM) were used to create DR configurations where virtual servers and data may need to be recovered at an alternate site. SRM allows some virtual server recovery operations to be automated, and makes operations like failover and DR testing easier. SRM does not have its own data mover though, and needs third party

SAN-based replication tools to address that issue. And again, any solution created using these products was specific to VMware and drove separate purchases to provide the same types of recovery capabilities for physical servers and/or other virtual server platforms that may be present in an environment.

Even when crafting a recovery solution that is specific to VMware, heterogeneity must be taken into account. VMware customers use a variety of different storage architectures in their environments, including SAN, DAS, and NAS. Popular applications commonly in use include Exchange, SQL, SharePoint, Oracle, various Windows, Linux, and/or Unix file systems, and SAP – each of which may have their own APIs for creating application-consistent recovery points.

Taking this discussion into account, VMware recovery requirements can be defined in this way:

- Data must be protected and recoverable using an approach that must be able to meet the strictest RPO and RTO requirements, both locally and remotely, and reliably supports application-consistent recovery options
- The recovery solution must have very low overhead in use on virtual servers, network bandwidth, and storage
- Applications must be protected and recoverable using an approach that can automate both failover and failback
- Configuration must be sufficiently flexible to support clustered and non-clustered virtual servers, multiple replication topologies (1 to 1, N into 1, 1 to N), and any guest operating systems (Windows, Linux, Unix), storage topologies (SAN, DAS, NAS), and application environments (messaging systems, databases, file systems, etc.)
- Data protection and application recovery operations must operate at the level of the individual virtual machine, not at the ESX Server level, to enable the most efficient operations
- All of these recovery capabilities must be available in a single recovery solution that covers backup, DR, and HA requirements to simplify recovery and keep costs to a minimum

While the primary topic of this discussion focuses on providing recovery for VMware environments, most customers will still have physical servers, and others may have other virtual server platforms, such as Microsoft Hyper-V and Citrix XenServer, in use as well. There is significant value in deploying a recovery solution which can accommodate each of these different environments with application-aware (not application-specific) recovery capabilities that has the same deployment model across all environments. Such a solution will be simpler, more cost-effective, and easier to manage across the entire environment than a patchwork collection of platform and application vendor specific tools.

Other Considerations

N Tier Environments

With the increasing deployment of SharePoint and SAP, N tier application environments are coming into wider use in enterprises today. These environments tend to have a back end database tier, an applications tier, and a front end or “presentation” tier. Customers looking to lower the physical infrastructure costs of these types of applications will often mix virtual and physical servers together. Virtual servers offer lower deployment costs for application services tiers, while the use of physical servers to host the database back end provides the performance and scalability that these types of environments demand.

N tier application environments configured in this way present two additional recovery challenges. First, creating a single, “synchronized” recovery point across tiers can be very difficult using conventional, point-in-time based “backup” approaches, resulting in additional complexity during recovery operations. And second, the mix of physical and virtual servers may demand at least two separate backup approaches, again complicating recovery and increasing costs. If N tier application environments are in use, these challenges are not well addressed by conventional backup vendors.

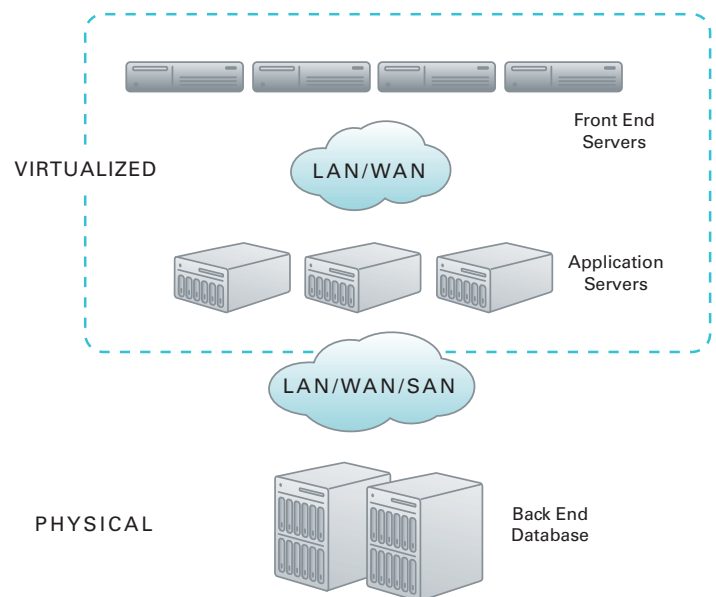


Figure 1. N tier application environments can often be more cost-effectively deployed if both physical and virtual servers are used as deployment platforms.

Virtual Storage Appliances

Particularly for smaller customers with lower performance requirements, the use of virtual storage appliances (VSAs) in protection and recovery operations can offer some cost-saving options. During a backup, most backup products create a discrete copy of the production data set on a separate server across a network, putting the backup load on a network that may be shared with other production servers. The use of a VSA can present some potential optimizations. Assuming that several virtual servers are configured using the same networked storage, a VSA can be brought up on one of the ESX Servers and designated as a “backup appliance.” When performing a backup, the data is transferred across the storage network, not the local area network (LAN), to effectively get it to the backup target (the backup appliance). Alternatively, a disk-based snapshot could be created on Server A (one of the production virtual servers), and then detached and mounted on the backup appliance using a VMware feature called “hot add.” Further processing (e.g. deduplication, compression, encryption, etc.) can then be done by the VSA before the backup is sent out across the LAN or WAN without necessarily impacting the performance of the other virtual servers running on that ESX Server.

The use of a virtual backup appliance can lower LAN bandwidth requirements and save on the hardware costs of a backup server, but it is important to understand the performance implications of VSAs if and when they are deployed. If a VSA is deployed on an ESX Server with other production virtual machines, they will all effectively be using the resources of a single physical server. It is important to configure that ESX Server appropriately to meet performance requirements. If the VSA requires a lot of CPU, memory, and/or storage resources, it may be better to configure it on its own separate physical server.

Individual Virtual Machine Granularity

For most enterprises, only a certain set of their servers are deemed critical enough to be protected using replication technologies for DR purposes. When configuring virtual server environments, all virtual servers on a given ESX Server may not be critical enough to warrant such protection. It is important to be able to configure protection and recovery solutions at the level of the individual virtual server (the *vmdk* file) rather than at the less granular level of the entire ESX Server. Replication consumes network bandwidth, and it would be most desirable to only have to replicate the data of critical virtual servers rather than all the virtual servers that may reside on a particular ESX Server. Being able to deploy replicated configurations with this level of granularity will conserve bandwidth and lower overall costs.

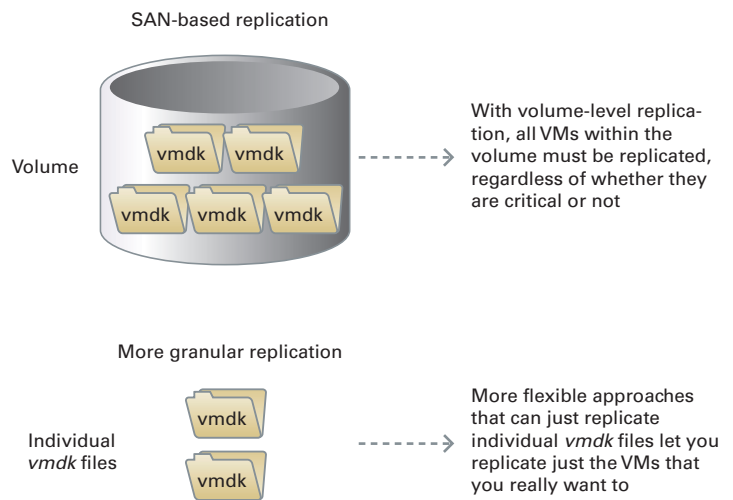


Figure 2. *If only 2 of the 6 vmdk files are really critical, then why replicate all of them? Just replicating the 2 critical ones results in an easier deployment that uses less bandwidth.*

While these comments on the use of replication technology apply specifically to data, granularity concerns also apply to the protection and recovery of application services. HA technologies in use in these environments should not only support automated failover at the level of individual virtual servers, but should also support automated failback. This gives administrators options to speed recovery and reduce recovery risk without having to also configure and manage failover and failback manually for non-critical virtual servers.

InMage Recovery Solutions for VMware Environments

InMage offers disk-based business application recovery solutions for both physical and virtual server environments. Our integrated recovery software supports both local and remote backup and DR, performing exceptionally well in heterogeneous IT infrastructures with mixed platforms and storage. These comprehensive solutions simplify recovery and lower costs by replacing multiple existing products across various platform and application environments with a single, centrally managed solution that addresses both data and application recovery.

InMage employs a number of next generation recovery technologies in the solution, including continuous data protection (CDP), application snapshot API integration, asynchronous replication over IP, application failover/failback implemented around a shared nothing cluster architecture, and integrated WAN optimization. CDP continuously captures changes to protected

applications in real time, making data recoverable the instant that it is created. Because of its granular data capture and ability to retroactively turn any previous application data state (within a defined retention window) into a disk-based recovery point, CDP supports near zero recovery point (RPO) and recovery time objectives (RTO) and provides the industry's best recovery solution from data corruption. Because of its very low overhead and low resource utilization, CDP effectively eliminates backup windows, making data protection operations transparent.

InMage is integrated with application vendor supported snapshot APIs that allow it to mark application-consistent recovery points in the CDP data stream with minimal impact to production applications. InMage uses standard, off-the-shelf application snapshot APIs like Windows Volume Shadowcopy Services (VSS), Oracle RMAN, and others to "bookmark" these types of points for quick and easy future reference. While these bookmarks are often turned into disk-based recovery points, these types of snapshots (called AppShots by InMage) can be mounted on any network attached server to meet a number of other operational needs such as test, development, analytics, reporting, data migration, etc.

To support long distance DR configurations without impacting production application environments, InMage uses asynchronous replication. InMage supports a number of different replication topologies without imposing any additional impact on production virtual servers, including 1 to 1, 1 to N, and N into 1. Integrated with InMage's replication are WAN optimization technologies, including compression, TCP optimization and I/O throttling (bandwidth shaping), to ensure that the minimum amount of data is transferred across the network to make information recoverable at target locations. Because InMage supports block-based CDP and replication, it is much more efficient than file-based approaches that require more network bandwidth.

InMage also supports automated failover and failback for virtual servers and the application services that run on them using a shared-nothing disk architecture. More flexible and less complex than shared disk failover products, InMage relies on its built-in replication technology to keep source and target data sets synchronized, and can support automated failover and failback to/from any target locations at the level of individual virtual machines. InMage includes productized, fully supported failover/failback templates built for specific applications that handle fault detection, selecting and mounting the recovery data set, re-starting an application service, and re-directing any network attached clients through "push" updates of AD and/or DNS routers.

For failback operations, InMage uses delta differencing technology to minimize the amount of data that needs to be sent back across the network to re-synchronize the primary site data set prior to failback. Automated recovery operations like those offered by InMage speed the recovery of application services while at the same time minimizing recovery risk.

InMage's ability to deploy application-aware solutions across the key platforms in use by enterprises today differentiates their offerings from alternatives. InMage covers physical servers (Windows, Linux, and Unix), virtual servers (VMware, Hyper-V, and Citrix), heterogeneous storage subsystems using different storage architectures (SAN, DAS, and NAS), and application environments including Exchange, SQL, SharePoint, Oracle, MySQL, Blackberry Enterprise Server, SAP, and any Windows, Linux, or Unix file systems.

InMage is a VMware Technology Alliance Partner.

The InMage Architecture

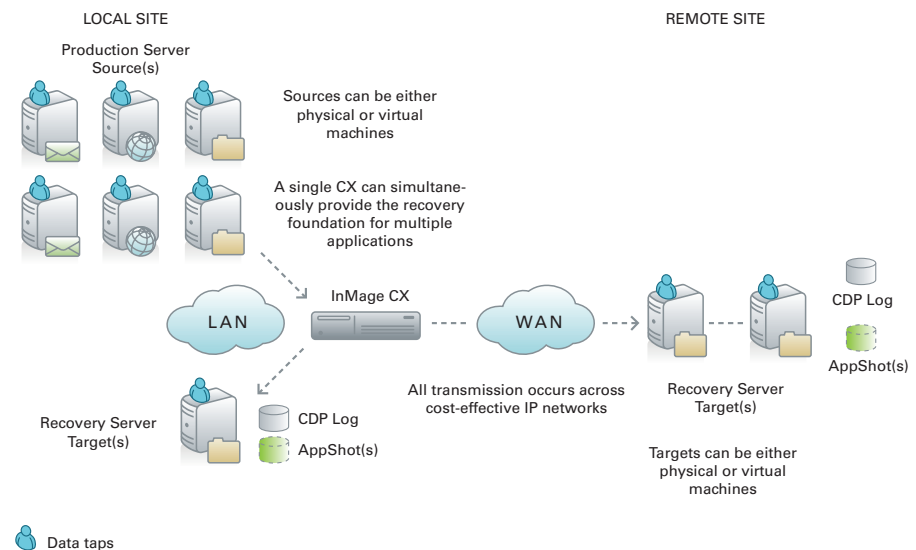


Figure 3. InMage captures data from production servers in real time, transmits it to an appliance (which can be local), and then from there can replicate it to one or more targets.

To capture the data necessary to rapidly and reliably restore data and/or application services running on virtual servers, InMage operates at two levels. At the first level a filter driver, which InMage calls a data tap, is deployed on each virtual server, right below the file system and above the volume manager. This data tap passes reads directly through to primary storage but mirrors any writes across the LAN to the InMage CX, an x86-based appliance running Linux. The CX can then WAN

optimize and/or encrypt the data stream prior to continuously and asynchronously replicating it across the WAN to one or more target locations. Recovery server targets that can reside locally, remotely, or both, host a CDP Log where the CDP data stream - time stamped, bookmarked, and write-ordered - is kept on disk. All data movement occurs across cost-effective, IP-based networks.

At the second level InMage uses the VMware *rcli* to interact directly with ESX Server at the hypervisor level to collect system configuration information, known as vmx data. Vmx data concerning any and/or all the virtual servers on a given ESX Server can be replicated to a single master target on the remote side, and InMage will use that data to help automate the process of setting up recovery server targets and recovering virtual servers and applications if and when necessary. Keep in mind that for any data capture, regardless of whether it comes from the data tap or through the *rcli*, InMage is using CDP to capture data continuously as it is created - we are not using changed block tracking because we don't need it - and we are doing this at the granularity of individual virtual servers, not ESX Servers. This gives administrators the ability to replicate, protect, and recover only the virtual servers that they need to, resulting in simpler configurations that use less network bandwidth and have fewer recovery steps.

Note that the data flow is from source server to CX appliance to target server, a hybrid approach that offers the advantages of both host and appliance based replication without any of the disadvantages. The data tap residing on the virtual server is a very lightweight component that performs write mirroring, while all the "heavy lifting" necessary to support InMage's rich recovery functionality is performed by the CX. The CX can be deployed on either a physical or a virtual server. This effectively off-loads the data protection operations from the virtual servers being protected, imposing only 2-3% overhead on each and virtually eliminating the concept of a "point in time" backup.

When a recovery point is required, that point is selected through InMage's graphical user interface (GUI) and immediately created and mounted on the target recovery server that hosts the CDP Log. Individual files or entire virtual servers can then be recovered across the network back to any production server (virtual and/or physical). AppShots are InMage's trade name for application-consistent recovery points, which are most often used by administrators because they support faster, more reliable recovery operations than crash-consistent recovery points. But InMage can retroactively generate any previous point in time, most of which are crash-consistent points, and these points may have value for root cause analysis and other administrative operations.

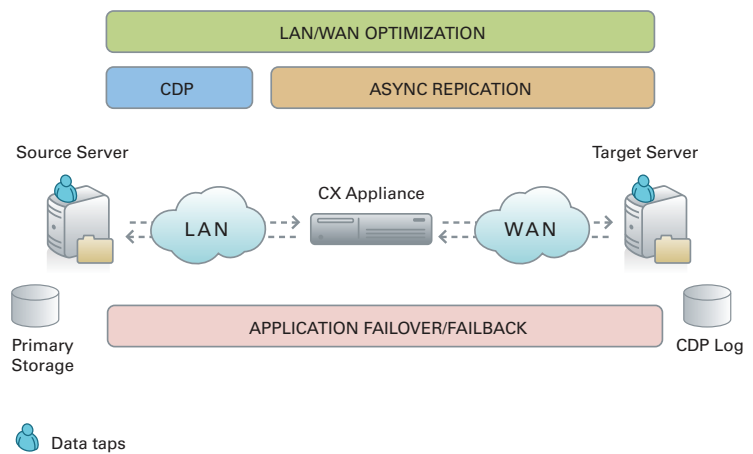


Figure 4. InMage uses a combination of next generation recovery technologies that eliminate backup windows, meet the strictest RPO/RTO requirements, and support excellent recovery reliability.

Critical InMage Recovery Capabilities for VMware

Relative to conventional backup products and VMware-specific recovery software, InMage offers some distinct advantages:

- Because InMage uses CDP, it eliminates backup windows and makes data recoverable as soon as it is created, not just once it's backed up; compare this to point in time-based backup approaches that impact production applications during the backup window and where your ability to meet RPOs is determined by how often you back up
- Because InMage has a filter driver in each guest OS image, we can interact directly with application snapshot APIs like VSS and RMAN, giving us the ability to track and make available reliable, application-consistent recovery points without imposing the overhead of conventional backup and snapshot agents; note also that marking application-consistent recovery points in the CDP data stream is much faster than actually creating a disk-based snapshot because we are not moving any data to insert the bookmarks, giving us much lower overhead than conventional snapshot approaches

- InMage is very efficient in its use of both network bandwidth and storage capacity for several reasons; first, we are capturing data at the block level (although our “book-marking” capability gives us the application-awareness that most block-based solutions do not have); second, after an initial synchronization we are capturing only change data, a fact which effectively removes most of the redundancies associated with backup data without having to incur the “processing” overhead of conventional data deduplication products; and third, we have integrated WAN optimization technologies which minimize the amount of data that has to be sent across networks to make information recoverable at target locations, and I/O throttling capabilities to ensure that InMage plays well with other applications with which it shares networks
- Because InMage can simultaneously protect both physical and virtual servers, we offer a single solution that can present synchronized recovery points across tiers for N tier application environments like SharePoint and SAP
- InMage’s ability to flexibly support P2P, P2V, V2P, and V2V configurations provides simplifying options for server-level recovery without having to deal with the complexities of bare metal restore
- Because InMage effectively offers a “4 in 1” recovery solution (backup, DR, HA, and WAN optimization) it replaces multiple existing products in use in not only VMware environments but also in Hyper-V and Citrix environments and on physical servers, greatly simplifying protection and recovery operations and resulting in an ROI of under a year in many cases

InMage supports VMware Infrastructure 3 as well as vSphere 4, and can accommodate virtual servers running a variety of different guest operating systems, including Windows, Linux, Solaris, AIX, and HP-UX.

Conclusion

As critical production applications increasingly deploy on VMware, recovery becomes a more pressing concern. Native VMware recovery tools are built around a “point in time” orientation to data protection that, for the more critical application environments, may lead to backup window, RPO, RTO, and/or recovery reliability issues. Separate tools like vHA, vFT, or SRM are needed to handle virtual server and application recovery.

Industry trends like very high data growth, evolving recovery requirements, and increasing heterogeneity with its associated costs and complexities are impacting VMware environments, and will be driving more enterprises using that technology to consider moving away from point in time based data protection to more continuous approaches.

InMage offers a comprehensive recovery solution, built around next generation recovery technologies like CDP, application snapshot API integration, asynchronous replication over IP, granular application failover and failback, and integrated WAN optimization, that simplifies protection and recovery operations in VMware environments by centralizing operations into a single tool that handles both data and application recovery, either locally or remotely. InMage is the only solution in the industry that employs the “continuous” protection approach necessary to eliminate backup windows, meet very stringent RPO and RTO requirements, and offer application-consistent recovery operations in a package that is low overhead enough for deployment in VMware environments. We offer the best application-consistent recovery solution for VMware, and can back that up with product functionality and customer references available only from InMage. And we are the only solution that can simultaneously protect physical and virtual servers, making us a perfect solution to provide recovery for N tier application environments that may mix these server types as well as any environment that has a mix of both physical and virtual servers (as most IT shops do these days).

InMage gives you the recovery capabilities that enable the deployment of the most mission-critical applications on VMware, providing better recovery using fewer tools and products than if you were to build snapshot/backup recovery using VMware’s own tools and any required extra cost third party products (such as SAN-based replication). And these same industry-leading recovery capabilities can be used to protect and recover data and applications on physical servers as well as other virtual server platforms like Hyper-V and Citrix – all using the same, simple, low overhead deployment model.