

Leveraging Next Generation Data Protection Technologies To Solve Your Recovery Problems

If you're like most companies, your data is growing at a very rapid rate, in some cases maybe even uncontrollably. And for a lot of companies, recovery requirements are also becoming more stringent for reasons that have nothing to do with data growth, being driven by either business or regulatory mandates. If you're using tape-based infrastructure for disaster recovery and/or local backup, there's a good chance that these two trends are starting to cause some real problems for you in providing the recovery capabilities you need to protect your business.

A variety of next generation data protection technologies promise to meet recovery requirements despite these two trends. These technologies are "next generation" in the sense that they leverage disk, rather than tape, to ensure recoverability. The key ones, along with the benefits they provide, are as follows:

Disk-based recovery. Leverages the improved performance (random access) and reliability advantages of disk vs tape to reduce backup windows and provide faster, more reliable recoveries.

Continuous data protection (CDP). Reduces data protection impacts on production environments to almost zero while supporting recovery with zero data loss and built-in protection against data corruption.

Heterogeneous, asynchronous replication. Enables disaster recovery solutions over any distance without any storage subsystem limitations or impacts on application performance.

WAN optimization. Leverages a number of different technologies to reduce the amount of data that has to be sent across WANs to support recovery from remote locations.

Recovery Automation. Improves the speed of recovery while simultaneously reducing the risk and cost associated with manual recovery operations.

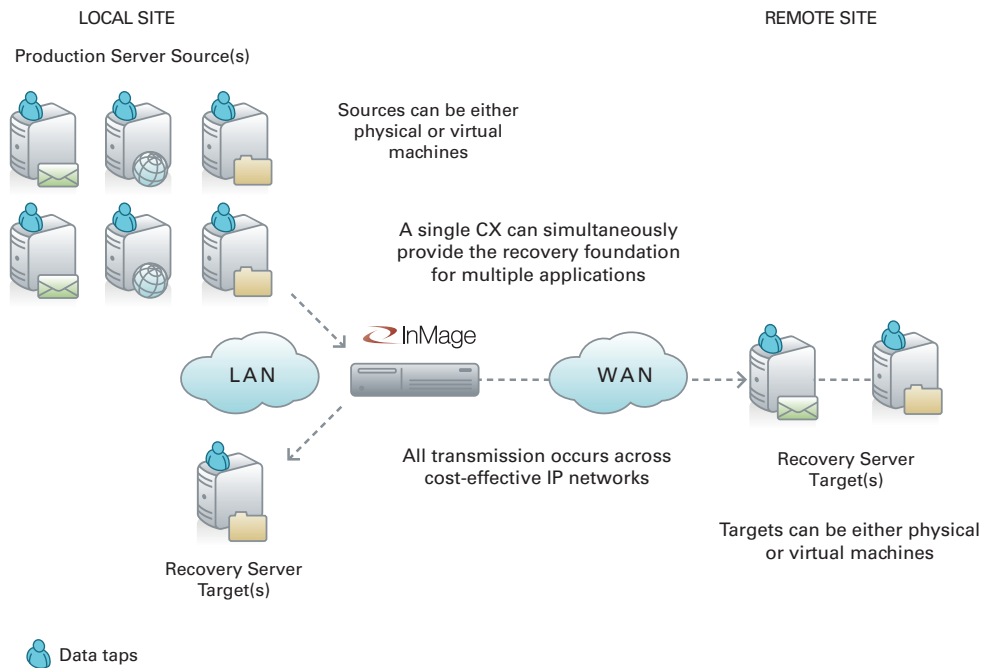
Simplifying Recovery

InMage combines all five of these next generation data protection technologies into a single solution that supports mixed physical and virtual machine (VMware, Citrix, Microsoft) environments, heterogeneous servers and storage, and any applications in configurations that provide disaster recovery, backup, application high availability, or any combination of the three. When used for recovery, InMage eliminates backup windows, improves recovery point (RPO) and recovery time objectives (RTO), and provides better recovery reliability for Windows, Linux, and Unix environments.

While all of our customers use InMage as a disaster recovery solution, over 40% of them simultaneously address local backup and recovery with it as well, and another 15% leverage the built-in application failover/failback to provide comprehensive coverage for both applications and data.

How It Works

InMage uses a unique hybrid recovery architecture that has minimal host impact, scales much better than host-based replication, and offers flexibility not supported by array-based replication products. Using InMage data taps (kernel-based filter drivers), data is captured from production servers across a LAN using CDP, and distributed to identified targets (either locally or remotely) using asynchronous replication. InMage works in SAN, NAS, and DAS storage environments, but does not require a SAN (all data movement occurs across cost-effective IP-based networks).



To perform a recovery, an administrator selects any previous point in time (up to and including the most recent point), and generates an "AppShot" (a disk-based, application-consistent recovery point). The AppShot is mounted on any network-attached server, and file or system-level recovery proceeds from there. Recovery kits provided by InMage automate application recovery for key enterprise applications, including Microsoft Exchange, SQL, and SharePoint, as well as Oracle, MySQL, Blackberry Enterprise Server, and SAP.

For more information on InMage solutions, you can visit us at www.inmage.com and download an introductory data sheet (InMage Scout Core Data Sheet) or an introductory white paper (InMage Scout Core White Paper).