HotLink

White Paper

Leveraging Public Cloud for Affordable Disaster Recovery & Business Continuity

By Edward Haletky
Principal Analyst
The Virtualization Practice

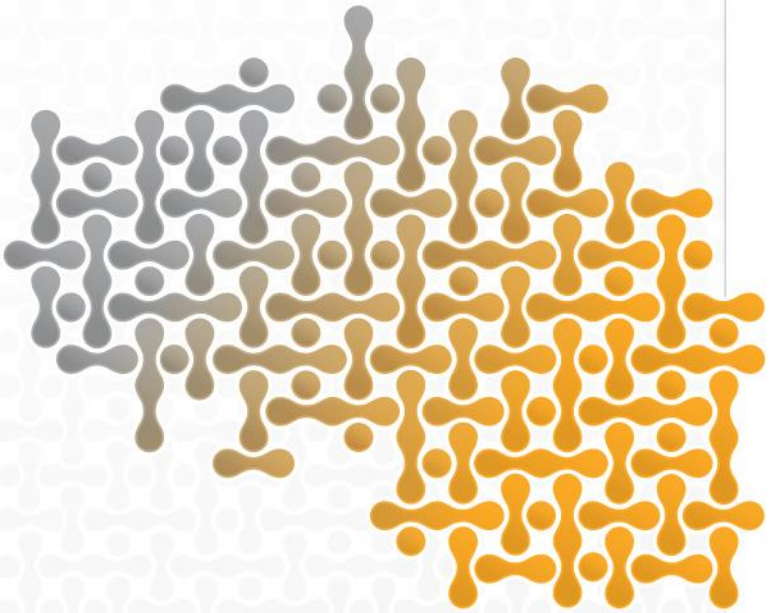

Table of Contents

Abstract

Every organization, whether large or small, needs to safeguard its IT assets in the event of some sort of interruption. Security breaches, networking issues, software problems, hardware failures and human errors are the most common culprits in such events, and basic backup is the typical solution given its affordability. Ideally, data protection would allow for quick and easy recovery of the data that was backed up, replicated, or stored within an archival unit. After all, without the ability to readily recover the data, backup is only so useful. The problem is the cost and complexity. Rapid recovery with limited data loss has been associated with a high price tag – the faster and more accurate the recovery, the greater the cost. As a result, the virtual machines in most organizations take days to recover from backup, assuming spare hardware is available. If new servers need to be purchased or if backups were not tested properly, the recovery period could stretch from days into a week or more.

Of course, there is a cost associated with each server that's down – not only in replacing the underlying hardware of the failed system, but also with loss of work. Engineering teams are a great example of the real cost of hardware failure. With tight deadlines and expensive personnel, the loss of a server and associated lack of productivity can have costs in the tens of thousands of dollars or more.

Productivity lost during downtime can be a big problem, so why don't organizations invest in more robust disaster recovery and business continuity solutions? Until now, the costs of these solutions were exorbitant, their configurations were inflexible and implementation was painful. They also required specialized skills, so only business-critical operations could justify the deployment. Most workloads do not meet that threshold of criticality, as they are either non-critical production systems or test and development workloads. Of all the VMs in operation, the vast majority are not considered mission-critical.

The real need for the overwhelming majority of deployed workloads is a solution that is dramatically better than backup, but cost-effective, simple to implement and provide continuous virtual machine protection. This solution will allow fully automated recovery, be easy for VMware system administrators to operate and test and enable manageability during recovery using the existing administration and management console – *VMware vCenter*. Best of all, the mirror site would be Amazon Web Services (AWS), so the price would be right and the scale unlimited.

This white paper is about a game-changing data protection technology from HotLink. This unique solution provides integrated VMware backup, replication, disaster recovery and business continuity in Amazon Web Services. Best of all, it's affordable for ordinary, everyday failures that happen in all IT environments. Read on to learn all about it.

State of the Data Protection Industry

The traditional data protection industry for VMware vSphere environments includes four main types of products at a wide range of price points and recovery timelines. License models span per VM, per de-duplicated TB, per CPU, or some combination. *All share one very important characteristic:* ***Restoration of operations requires spare VMware vSphere capacity (compute resources, memory, storage and networking), whether on-premise, in the cloud, at a service provider or a hosting facility.***

Let's review the traditional data protection options:

1) **Backup** – This is the most economical and pervasive of all data protection approaches. Traditional backup tools will back up VMs to a target store, either on premise or in the cloud. The data store needs to be carefully managed, as storage infrastructure is often expensive and capacity may be limited, particularly on premise. To that end, most backup tools will de-duplicate in order to make more efficient use of the storage. Since most solutions do not self-test, virtual machines may not be restorable if inadequately tested manually. Also, backup tools have their own management consoles that require skills in how to use. Restoring from backup typically takes many hours or days depending on the environment, quality of backups, and specific tool used.

2) **Replication** – These tools replicate data from one virtualization cluster to another, and the second cluster can be within an alternate data center, hot site, rack of computers or even a cloud. Replication requires that the secondary site is virtually *identical* to the original. For hardware-level replication, the storage hardware must match, doubling the cost. Compute, memory and network should also match for best results, and you will likely need to maintain an additional VMware vCenter Server at the secondary environment. Finally, the settings of the network configuration need to be automated, which is outside the scope of the replication products. All these factors require considerable preparation, ongoing testing and consistent management to ensure both sides of the replication environment are perfectly in sync. The restoration process can range from 15 minutes to hours, if not well-tested, well-prepared or working as planned.

3) **Near-continuous data protection –** These solutions can be considered replication on steroids, so the cost increases substantially. It can take place outside or within the virtual machine. Once again, a mirror site is required at a remote location, cloud or otherwise, so an investment in duplicate hardware and software or cloud-based services is required. As with standard replication tools, the target site must be carefully constructed, managed and tested. The separate management console(s) required in this environment also require specialized skills and training. In this class of data protection, the restoration process, assuming it's problem-free, is typically under one hour.

4) **Continuous Data Protection –** Typically reserved for mission-critical applications, this option is the most expensive by orders of magnitude. Besides inheriting the mirror site redundancy requirements of the near-continuous data protection environment, very low-latency networking and hardware is mandatory. Given the low tolerance for downtime and data loss in this solution category, robust management, monitoring, testing, and other requirements well outside the scope of this paper are also pre-requisites. This class of data protection requires recovery time in seconds.

All the traditional options for data protection share the following limitations that ultimately result in a *significant total cost to recover* and readily continue operations in the event of a failure:

- Need for standby capacity (storage, network, compute, memory)
- Need for skills and knowledge of data protection software management tools and maintenance
- Need to manage and test standby or hot-ready capacity (even in the cloud)
- Need to maintain and test networking and other necessary configurations for standby locations
- Need to commit personnel to build, test and maintain overall data protection environment
- Need for consistent VMware vSphere hypervisors for recovery, whether on or off premise

If you have proper capacity, configurations and good backups, then restoration should work. If there is ready capacity, configuration and good replications, then restoration should work. However, as the number of workloads increase, the amount of on-premise ready spare capacity decreases. For many, spare capacity is not in the budget at all. And most IT managers will admit, at least privately, that testing is not at the level needed because it's time consuming and the staff is stretched too thin. In combination, the lack of spare capacity and time-consuming testing render traditional data protection insufficient for a majority of VMware installations today.

Next Generation Data Protection

The next generation of data protection solutions not only needs to provide for backup and replication but also disaster recovery ***and*** business continuity for *every type* of VMware workload, not just business-critical ones. Key to making this possible is the ability to leverage resources in an elastic cloud, such as Amazon Web Services, that provides consumption-based pricing plus near infinite capacity. The next generation can't just provide backup to public cloud; it must also provide rapid recovery in the cloud. Of course, a must-have is easy migration back on premise once new hardware is in place or other resources are made available for restoration.

Consumption-based AWS pricing, which is pennies on the dollar compared to on-premise spare capacity, will dramatically lower the overall cost of DR/BC, so the capabilities will become viable for day-to-day workloads in engineering, development, test, internal systems and other non-critical applications of virtual infrastructure. These are the very users that only have budgets for basic backup and replication tools today, and they represent a majority of the VMware install base. Whether the failure is switching, blade chassis, systems, storage or some other hardware crucial to running the workloads, rapid recovery in the cloud will enable users to stay productive during a hardware outage while new equipment is installed and configured.

To be usable by the mainstream users, the next generation of data protection needs to provide comprehensive disaster recovery and business continuity that are easy to implement without professional services, simple to test without time-consuming processes, intuitive to use without specialized training, and easily manageable without yet another console. It needs to leverage the existing VMware vCenter administration and management infrastructure.

A summary of the core characteristics of next-generation data protection for VMware users include:

- Low cost and affordable data protection
- Business continuity, disaster recovery and backup capabilities
- Business continuity for ALL workloads, not just mission-critical ones,
- Near infinite spare capacity within an elastic cloud like AWS
- Recovery managed from within VMware vCenter
- Utilizes existing on-premise hardware
- No need for like-to-like backup, restoration or recovery

While today's data protection market has many vendors, tools and approaches, only one solution meets these next-generation data protection requirements – *HotLink DR Express™*.

HotLink DR Express

The next-generation HotLink DR Express is the only solution to extend VMware vCenter administration and management to include comprehensive data protection in Amazon Web Services. Not only does it backup and replicate VMware virtual machines and differentials in Amazon S3, but it enables VMware workloads to be restored in *minutes* in Amazon EC2 if an outage happens. The recovered AWS workloads are fully manageable by VMware vCenter alongside on-premise vSphere virtual machines. Most importantly, HotLink DR Express provides all these capabilities at prices any organization can afford.

HotLink DR Express is a plug-in to VMware vCenter that enables native administration and management of Amazon EC2 resources. AWS accounts and instances are treated just like vSphere hosts and VMs in VMware vCenter, so the user experience is completely intuitive. This is possible as a result of the patented HotLink Transformation Engine™, which handles the platform transformation, so VMware vCenter can natively interoperate with the Xen-based Amazon platform. Native integration with VMware vCenter Server enables HotLink DR Express to leverage VMware snapshots, virtual disk cloning, change block tracking, and other VMware vCenter administration and management capabilities for the DR/BC site in AWS. Using VMware vCenter, administrators select the VMs, templates or services to be protected along with the frequency of updates, with up to 500 restore points per VM. Seeding protected VMs and ongoing differentials to AWS is automatic, and no other processes or specialized tools are required.

The following diagram shows the flow of data with HotLink DR Express. Data moves from VMware vSphere as deltas with up to 500 restore points per VM. This data is stored in Amazon S3 in a compressed and encrypted format native to vSphere. If an on-premise failure happens and a restore is triggered in VMware vCenter, HotLink DR Express collapses the deltas and transforms the VMware vSphere virtual disk image and configuration files into an Amazon AMI. The Instance is powered on and available for use in minutes. While running in Amazon EC2, the workload is fully manageable by VMware vCenter, side-by-side with on-premise vSphere VMs.

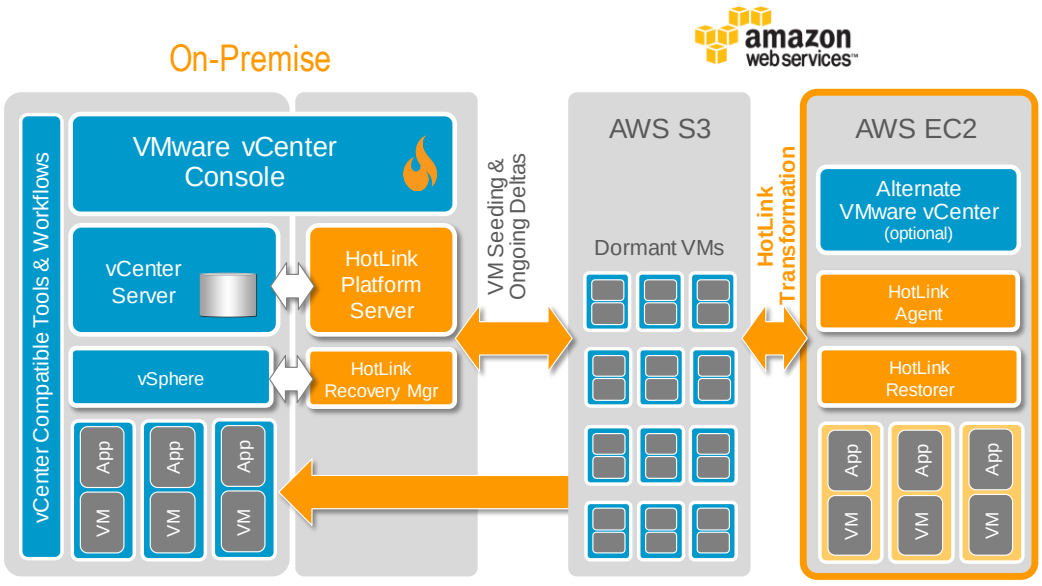

In addition, the integrated and automated HotLink workload transformation provides automatic conversion between native image formats so Windows and Linux VMDKs can be converted to Amazon Machine Instances (AMIs) and vice versa, a capability critical to providing seamless, bi-directional interoperability between VMware vCenter and AWS.

HotLink DR Express provides hybrid software-defined networking (SDN), so on- and off-premise networking can be seamlessly and holistically managed and pre-configured with VMware vCenter as the integration hub (see below diagram). By automatically discovering, transforming and mirroring on-premise network configurations in AWS, HotLink DR Express removes the complexity and labor-intensive nature of managing hybrid DR networking, maintains networking consistency across application tiers in a recovery, and enables easy DR testing without disturbing production operations. The most important benefit is that with pre-configured, automated network provisioning in AWS, VMware workloads can be accurately and automatically restored in Amazon EC2 in minutes, so operations are minimally disrupted following an on-premise failure.

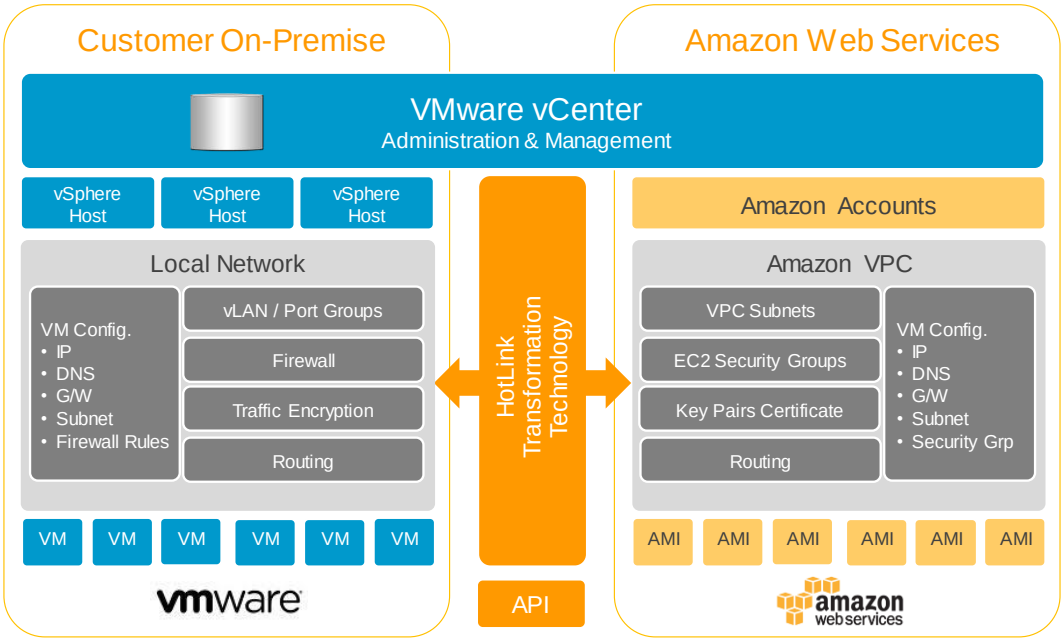

Once the on-premise or alternate VMware site has available capacity, the workload can be automatically cloned back at any restore point in Amazon S3, or the running Amazon EC2 workload can be cold migrated, cold cloned, or hot cloned. This can happen whether the EC2 Instance is running or powered off (hence the hot or cold options). This AMI to VMDK conversion is an integrated and automated function of HotLink DR Express.

Following is a brief look at the key functionality of HotLink DR Express, providing all the features required of the next generation of data protection, but at the price of basic backup solutions.

- Single point of comprehensive data protection management using VMware vCenter
  - Extend VM administration & management capabilities to Replication/DR/BC site
  - Build recovery plans in PowerCLI, Orchestrator or other vCenter compatible tool
- Automated Replication/DR/BC site setup in Amazon
  - Configure RPO per service / VM with vCenter
  - Replicate on-premise resources & map on-premise services
- Seamless cloud Replication/DR/BC site maintenance
  - Replicate VMs & templates in AWS
  - Synch only deltas as on-premise VMs change
  - Monitor site, validate status & configure alerts using vCenter
- Integrate rapid cloud replication/DR/BC site testing
  - Provision test environment per IT service
  - Integrate & automate continuous service testing
- Restore operations in AWS in minutes
  - Recover in any AWS region, multiple AWS regions, or alternate on-premise hosts
  - Automate builds of AWS Instances or on-premise VMs at any restore point
- Migrate back on-premise when ready
  - Deploy current state or original restore point when on-premise site is available

Ultimately, data protection is about recovering from inevitable infrastructure failures and keeping operations running, whatever the cause. With HotLink DR Express, every organization, large or small, has the benefit of including all VMware workloads in a comprehensive data protection at a low cost that is practical for everyday workloads — easily managed by the VMware vCenter infrastructure already installed.

Conclusion

The next generation of data protection must provide easy, well-integrated and intuitive backup, replication, disaster recovery and business continuity — *but at affordable prices*. This is exactly what makes HotLink technology a unique and compelling solution for the full spectrum of workloads. In fact, only with HotLink can VMware shops:

- Extend VMware vCenter capabilities to include backup/replication/DR/BC
- Restore operations in AWS in minutes
- Automate replication/DR/BC site setup & maintenance
- Test & validate the replication/DR/BC site easily
- Manage the AWS replication/DR/BC site with VMware vCenter
- Recover VMs on-premise automatically

HotLink transforms data protection into exactly what it should be: (a) cost-effective, (b) feature-rich, (c) well integrated into daily operations, (d) intuitive to use, (e) easy to text, and (f) able to eliminate downtime in minutes.

Check out the HotLink data protection products and services to see if they might provide the transformation your organization needs to protect against the next inevitable failure in your IT infrastructure.

About the Author

Edward Haletky is the author of *VMware vSphere and Virtual Infrastructure Security: Securing the Virtual Environment* as well as *VMware ESX and ESXi in the Enterprise: Planning Deployment of Virtualization Servers*. Edward is a principal analyst at The Virtualization Practice and owns AstroArch Consulting, Inc., providing virtualization, security, network consulting and development.

HotLink Corporation
3130 De La Cruz, Suite 211
Santa Clara, CA 95054
(408)463-6130
www.hotlink.com
info@hotlink.com