

WHITE PAPER

Optimized Backup and Recovery for VMware Infrastructure with EMC Avamar



EMC²

Contents

Introduction.....	1
VMware Infrastructure Overview	1
VMware Consolidated Backup	2
EMC Avamar Overview	3
Solution Overview	3
Backup Based on VMware Consolidated Backup	3
Guest-Based Backup	5
ESX Server Service Console-Based Backup	6
Appendix.....	8

Introduction

VMware Infrastructure is industry's most widely deployed virtualization solution. Virtual machines deployed in the datacenter must be protected against failure. Extending data protection to virtual machines is thus an important function. In the virtualized environment provided by VMware Infrastructure, there are many ways to improve the convenience and reliability of data protection, each with its particular advantages and challenges.

EMC Avamar is a data-protection solution for virtual infrastructure that uses unique deduplication technology to back up virtual machines efficiently without overly taxing the infrastructure. This paper discusses details of the Avamar solution and various ways of providing data protection.

VMware Infrastructure Overview

VMware Infrastructure is the industry's first full infrastructure virtualization suite that allows enterprises and small businesses alike to transform, manage, and optimize their IT infrastructure through virtualization. VMware Infrastructure delivers comprehensive virtualization, management, resource optimization, application availability, and operational automation capabilities in an integrated offering.

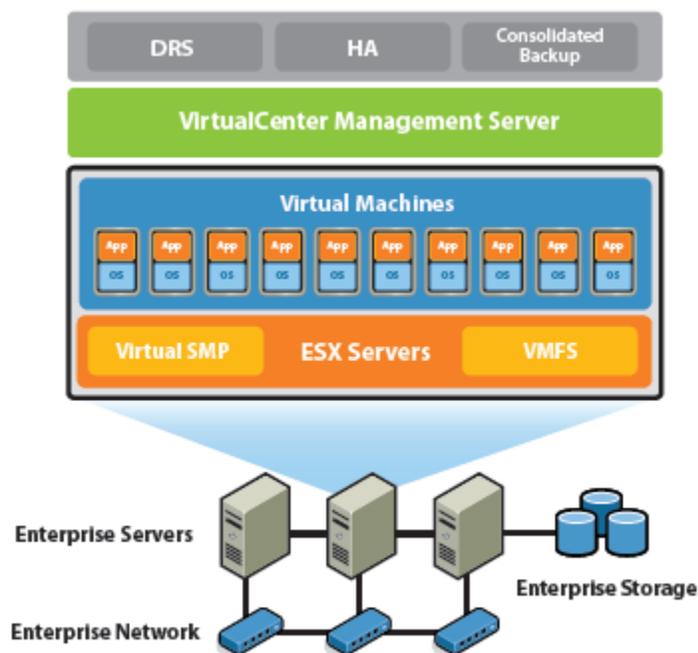


Figure 1: VMware Infrastructure

VMware Infrastructure includes the following major components:

- VMware ESX Server — A production-proven virtualization layer run on physical servers that abstracts processor, memory, storage, and networking resources to be provisioned to multiple virtual machines.
- VMware Virtual Machine File System (VMware VMFS) — A high-performance cluster file system for virtual machines.
- VirtualCenter Management Server — The central point for configuring, provisioning, and managing virtualized IT infrastructure.

- Virtual Infrastructure Client (VI Client) — An interface that allows administrators and users to connect remotely to the VirtualCenter Management Server or individual ESX Server installations from any Windows PC.
- VMware VMotion™ — Enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity.
- VMware Consolidated Backup— Provides an easy to use, centralized facility for agent-free backup of virtual machines. It simplifies backup administration and reduces the load on ESX Server installations.

VMware Consolidated Backup

VMware Consolidated Backup takes the backup load off the ESX Server host, eliminates the backup window, removes backup traffic from the LAN, and eliminates the need to run backup agents inside virtual machines to perform file-level backups of virtual machine data.

Consolidated Backup uses VMware Tools to quiesce the file system inside the virtual machine, ensuring that when the snapshot is taken, all pending data changes have been written to disk so the snapshot contains consistent data. Consolidated Backup also facilitates running scripts before and after the backup, so you can freeze and quiesce applications, then unquiesce them after the snapshot is taken.

Once the snapshot is taken, a separate physical machine — the backup proxy — mounts the base disk as if it were a locally attached file system so a backup agent running on the proxy can read and back up the files using the same features the agent makes available for backing up physical drives, as shown in Figure 2.

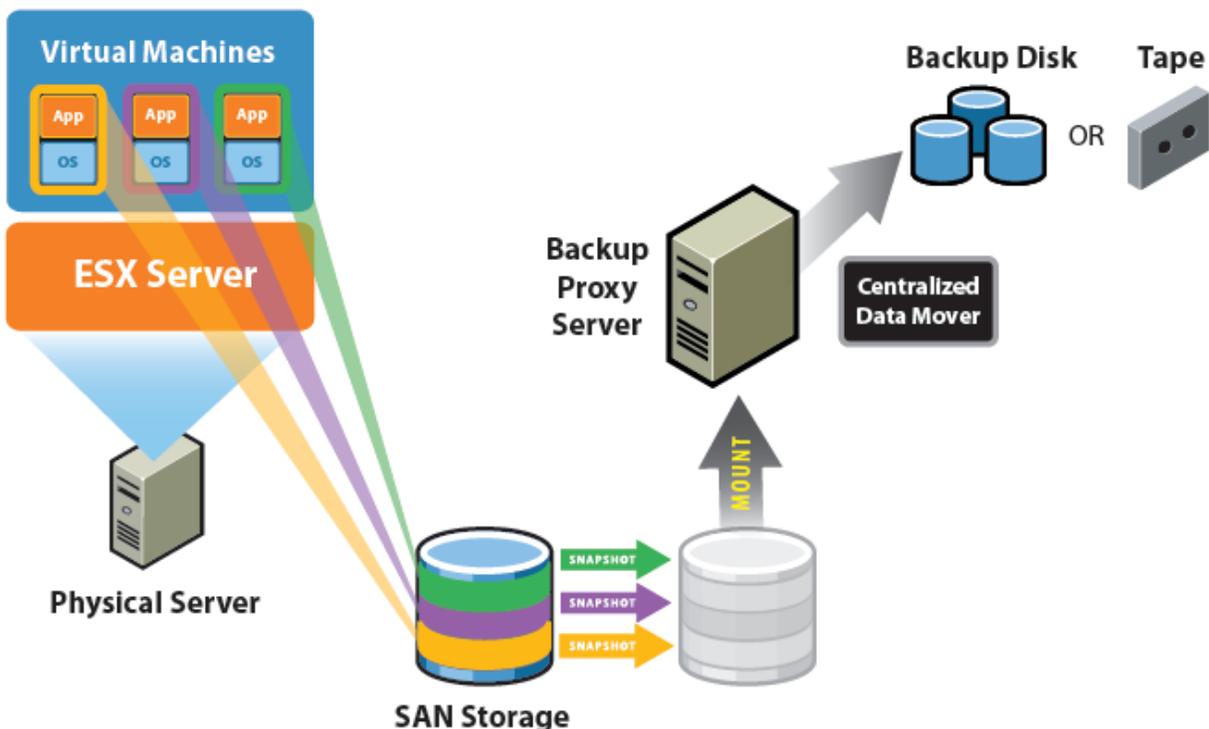


Figure 2: How VMware Consolidated Backup works

EMC Avamar Overview

EMC Avamar software is backup and recovery software — with an important difference. Avamar's global data deduplication technology eliminates the unnecessary transmission over the network and storage of redundant backup data. This deduplication slows the pace of data growth both in core data centers and at remote offices. Avamar is especially good in areas where traditional backup solutions struggle — virtual machines, remote offices, and large, LAN-attached file servers.

Traditional backup solutions require a rotational schedule of full and incremental backups, which move a significant amount of redundant data week over week. Because of the unnecessary data movement, enterprises are often faced with backup windows that roll into production hours, network constraints, and too much storage under management. In VMware Infrastructure environments, server consolidation can mean overlapping backup windows and heavy impact on hardware resources.

Over the last decade, disk storage has been used to augment traditional backup approaches, but disk solutions that are designed to replace tape libraries and media solve only a fraction of the data protection challenges faced by enterprises.

Avamar backup software solves traditional backup challenges by reducing the size of backup data at the source. Avamar does this using patented global data deduplication technology that stores only a single copy of sub-file data segments across all sites and servers.

Avamar software quickly and efficiently protects VMware Infrastructure environments by reducing the size of backup data within and across virtual machines — using agents in the virtual machines, on the VMware Consolidated Backup server, or in the ESX Server service console. For virtual machine backups, Avamar eliminates traditional backup bottlenecks caused by the large amount of data that must pass through the same set of shared resources — the physical server's CPU, Ethernet adapter, memory, and disk storage. Avamar reduces the traditional backup load — up to 200 percent weekly — to as little as 2 percent weekly, dramatically reducing backup times and resource utilization. Avamar also brings benefits whether an organization uses VMware Consolidated Backup or does backups from the ESX Server service console, as outlined below. Unlike traditional backup solutions, Avamar can deduplicate the data stored in virtual disks (.vmdk files).

Solution Overview

Different data center environments favor different solutions. EMC Avamar in a VMware Infrastructure environment offers you the flexibility of implementing your data protection solution in any one of three ways:

- **VMware Consolidated Backup:** Avamar Agent is installed on the VCB proxy server
- **Guest-based backup:** Avamar Agent is installed inside the virtual machine
- **Service console-based backup:** Avamar Agent is installed in the service console

Backup Based on VMware Consolidated Backup

VMware Consolidated Backup enables LAN-free backup and offloads backup workload to the Consolidated Backup proxy Server. The Consolidated Backup proxy can mount a Windows virtual machine's .vmdk files to provide a file system-level view of the virtual machine. Using the Avamar Agent to back up the mounted virtual machine disks, Avamar provides data deduplication at both the file level and the .vmdk level.

Consolidated Backup consists of a set of utilities and scripts that work in conjunction with Avamar Agent software and AVIM (Avamar VCB Interoperability Module). The Avamar Agent and the Consolidated Backup interoperability module run on the Consolidated Backup proxy to provide the backup services. Actual backup of the virtual machine happens on the backup proxy server. One backup server can provide backup services to many ESX Server hosts as long as all the machines share same SAN (storage area network).

The Avamar VCB Interoperability Module is a set of scripts that leverage Consolidated Backup scripts to create snapshots and to mount and unmount the snapshots. — point-in-time copies of the running virtual machines' drives. When the Avamar backup server initiates a backup according to the schedule and policy you specify, Avamar Agent on the proxy server initiates the backup action.

As the first step in backing up the virtual machines, AVIM scripts are called before the actual backup can happen. You can custom configure the scripts.

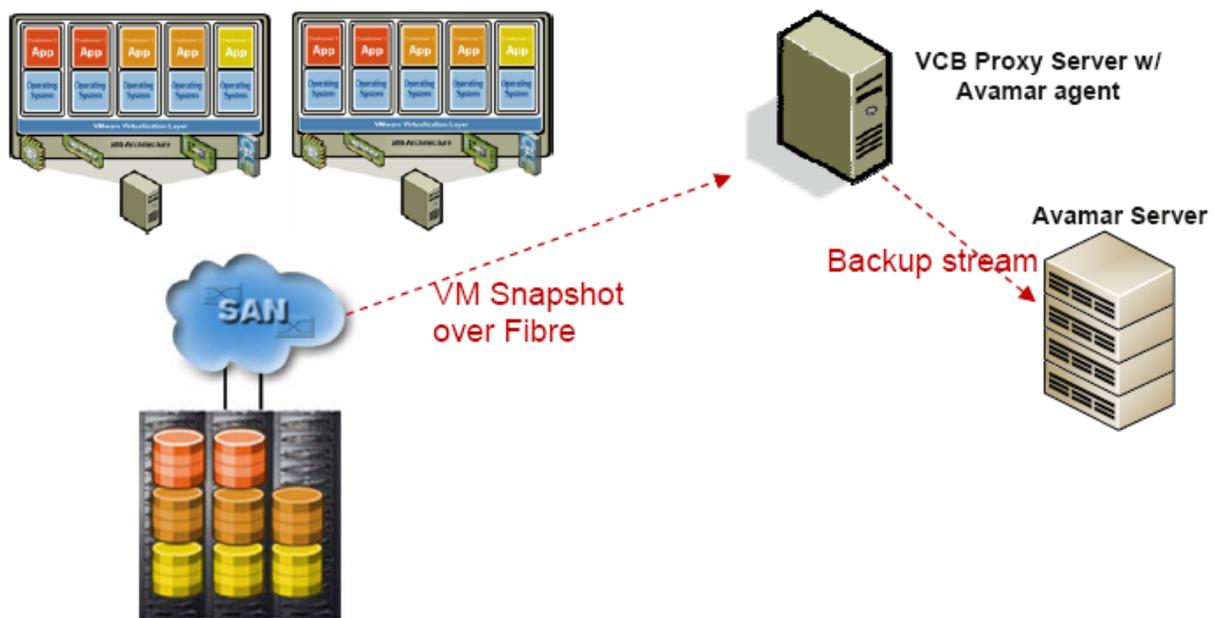


Figure 3: Avamar and Consolidated Backup working together

Advantages of using VMware Consolidated Backup are:

- Eliminates the need to schedule backup window downtime in order to perform nondisruptive backups
- Makes it possible to remove backup traffic and load from the IP network by moving data over the storage network
- Removes the processing load of backup from the ESX Server hosts
- Eliminates the need to manage backup agents in each virtual machine for most scenarios

Avamar backup based on VMware Consolidated Backup offers you two choices when it comes to restoring the virtual machine. You can either restore the entire virtual machine directly to the destination ESX Server host or to the proxy server.

- Restore directly to the ESX Server host
 - Requires Avamar Agent installed on the ESX Server host
- Restore to the VCB proxy
 - Manually copy the virtual machine to the ESX Server host
 - Run the Consolidated Backup restore command with proper options

For guests running Windows operating systems, you can restore individual files directly to the virtual machines. This operation requires Avamar Agent inside the virtual machine. For Linux guest, currently only image-level backup is supported.

Guest-Based Backup

Guest-level backup involves installing Avamar Agent inside each virtual machine

Backup configuration for this method is no different from that for a physical server. Usually no scripting is needed for this type of backup. Configuration beyond basic client setup might be needed to support a specific application, such as Microsoft SQL Server, Oracle, or Microsoft Exchange. The main advantages of this procedure are as follows:

- Highest level of data deduplication.
- Support for backup of applications inside the virtual machines
- Support for partial or file level restores
- Identical backup methods for physical and virtual machines
- No requirement for advanced scripting or VMware software knowledge
- Unchanged day-to-day procedures for backing up

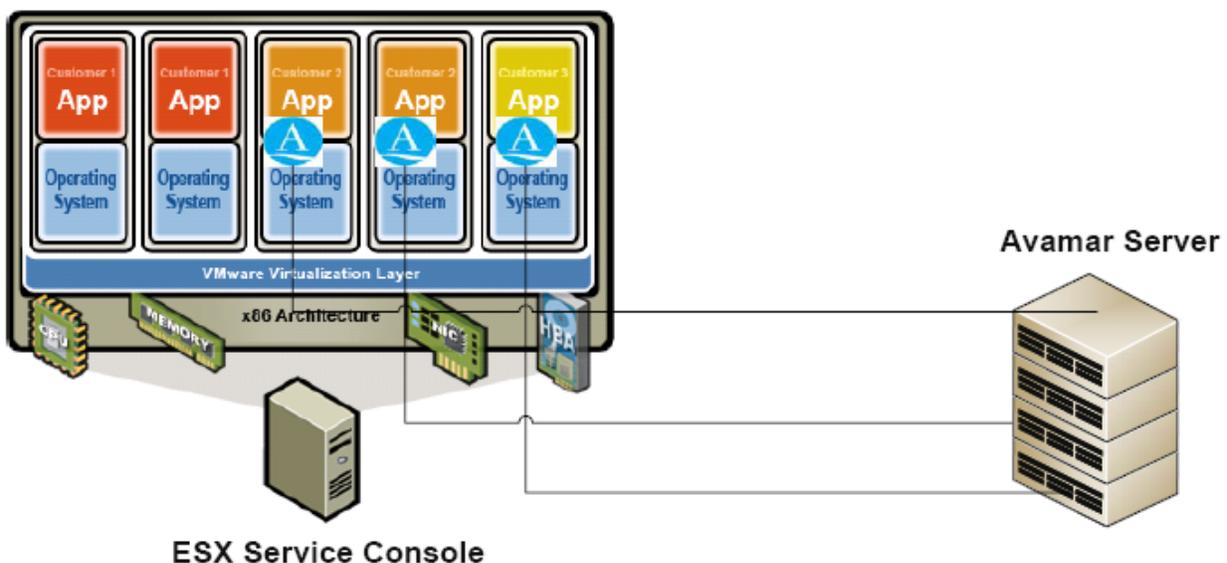


Figure 4: Guest- based backup

On the other hand, this is the most resource-intensive backup method. It also does not take advantage of such virtual machine characteristics as encapsulation for doing efficient backups.

Restoring virtual machines in this configuration is easily accomplished. Figure 5 shows the method for restoring a full virtual machine image.

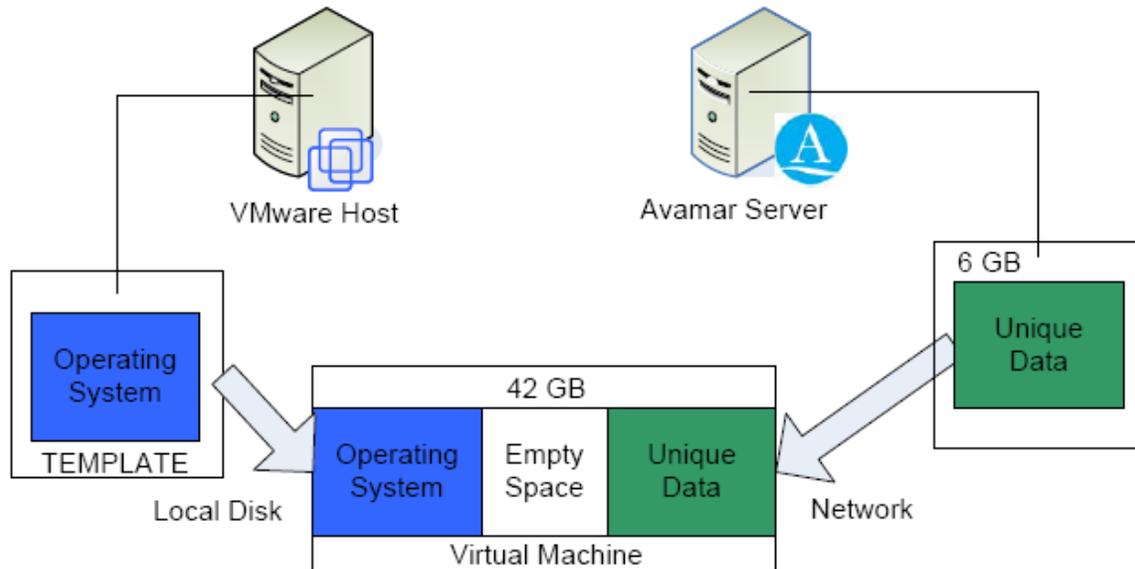


Figure 5: Restoring an entire virtual machine

To perform a full system restore (sometimes called a bare-metal restore), the system must create a virtual machine with operating system and Avamar Agent installed. This step is most easily achieved either by using a template or by using a ready-made operating system image. To restore an entire file system on the server, take the following steps:

1. Deploy a new virtual machine from a template or image.
2. Power on the virtual machine and register it with the Avamar server.
3. Perform a redirected restore to the new virtual machine.

ESX Server Service Console-Based Backup

ESX Server service console-based backup is the third option available to Avamar customers. It involves installing Avamar Agent on the service console of each ESX Server host that needs protection. This backup method combines the advantages of efficient management and image-level deduplication. Using the service console-based backup method, you can choose to back up the virtual machine either online or offline. In either case, you start the backup on the production ESX Server host and resources from the ESX Server host are used for moving the data to the Avamar server.

Another option when using this method is to back up virtual machines after suspending them or powering them off. In this case, you perform a simple backup of the `.vmdk` files for each virtual machine.

The ESX Server service console-based backup method is suitable for enterprises that can afford some downtime in the production environment.

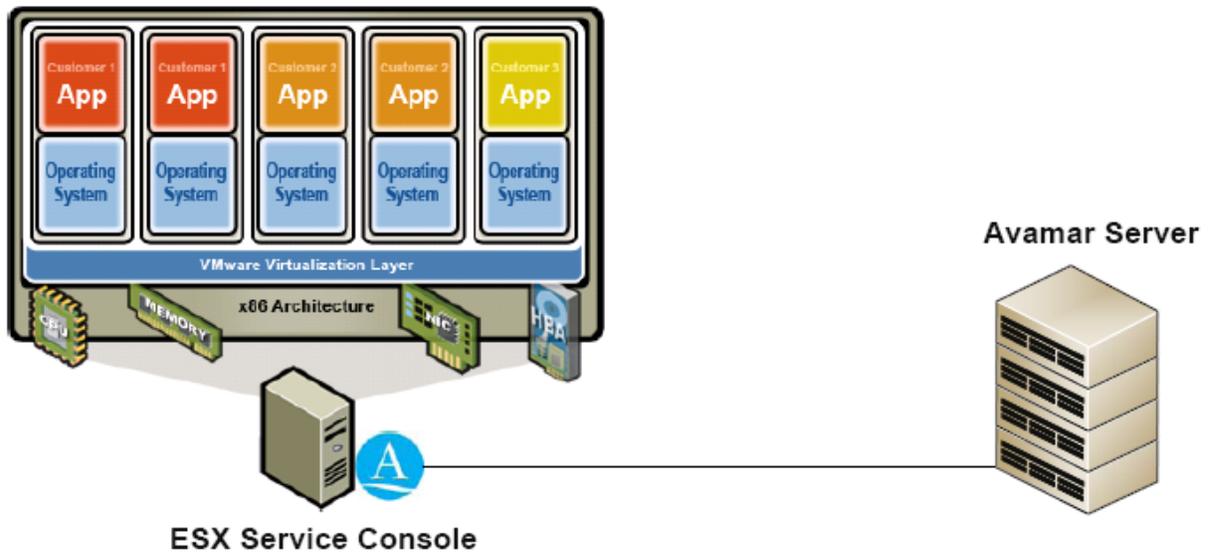


Figure 6: Service console-based backup

Advantages of this method of backup are:

- Less management overhead because backup agents are not needed inside the virtual machines
- No backup proxy server or SAN required because backup is done on the ESX Server host itself
- Restore goes directly to the ESX Server host

Restoring a virtual machine can be accomplished in one of two ways when you use this backup method. Two steps are needed to restore a virtual machine: restore the `.vmdk` files to the ESX Server host's VMFS file system, then register the virtual machine with the Avamar server.

If you use manual restore without the benefit of Avamar Agent, you have the option of manually registering the virtual machine.

Appendix

Supported guest operating systems for Avamar are shown in the following table:

Guest Operating System	Avamar Agent	Revision Level
Red Hat Linux 7.3	RHEL 3	3.7.0
Red Hat Linux 8.0	RHEL 3	3.7.0
Red Hat Linux 9.0	RHEL 3	3.7.0
Red Hat Enterprise Linux 3 (64-bit)	RHEL 3	3.7.0
Red Hat Enterprise Linux 4 (64-bit)	RHEL 4	3.7.0
SUSE Linux Enterprise Server 8	SLES 8	3.7.0
SUSE Linux Enterprise Server 9	SLES 9	3.7.0
Windows NT	Windows	3.7.0
Windows 2000	Windows	3.7.0
Windows XP	Windows	3.7.0
Windows Server 2003	Windows	3.7.0
Windows Server 2003 64	Windows	3.7.0

Revision: 20070727 Item: WP-024-ISV-01-01



VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806, 6,944,699, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, and 7,222,221; patents pending.
VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

