# Virtual Server System and Data Protection, Recovery and Availability with CA ARCserve® r16

# Table of Contents

# Introduction

Organizations are increasingly turning to server virtualization as a way to reduce IT costs and improve the manageability of their server infrastructure. However, virtualization can potentially mean a large, single point of failure. If the host system itself or its storage fails, the consequences can be catastrophic and impact multiple systems at the same time. Therefore, ensuring the protection, recovery and availability of your virtual server environment is critical. Very few organizations are 100 percent virtualized, and most organizations need to protect and recover a mixed environment of physical and virtual servers. Ideally, therefore, there should be a single solution that meets the needs of both physical and virtual environments.

However, virtual server environments do enable a flexible and easily deployed disaster recovery strategy. With the right tools, it can be much faster and easier to recover your systems, applications and data to a virtualized environment than having to build multiple physical servers. This paper provides a technical overview of how the CA ARCserve® Family of Products provides protection, recovery and availability of your systems, applications and data for both physical and virtual server environments.

# Key challenges for protecting virtual servers

As with any backup and protection strategy, it is important to ensure that the technologies used are appropriate and address the challenges in your environment. For protecting virtual servers, there are several specific challenges that any solution must be able to meet. These include selecting the right level of protection, providing for granular recoveries, managing large data volumes and backup performance and automating as much of the backup and protection process as possible.

For virtual servers, it is essential to select an appropriate protection level that meets your recovery objectives:

- **Host-based backups**. These protect the whole host server, and all virtual machines (VMs) running on that host, but may not be able to provide granular restores of individual VMs, or of applications or files on those VMs.
- **Guest-level backups**. These protect individual VMs, and provide granular restore of VM data, but depending on the application, may not be able to provide application-specific protection.
- **Application-level backups**. These protect individual applications running on the VM, such as Microsoft® Exchange, Microsoft SharePoint® and Internet Information Services (IIS). They offer granular restores of application-specific data, such as Exchange mailboxes or individual messages, or SharePoint folders.
- **Replication**. This complements periodic backups and snapshots by capturing every change made to data, files and databases, and eliminates the storage device as a single point of failure. Replication may be performed off site for disaster recovery. It typically includes a data rewind or continuous data protection (CDP) capability to rewind back to a known good point in time prior to the data loss or damage.
- **High availability**. This provides the best system and application recovery time by monitoring systems and applications and automatically failing over to another server or virtual machine to prevent an unplanned outage.

It is no longer enough to attempt to protect virtualization environments by simply installing data and system protection agents within the guest operating system, as if the guest VM were no different from a physical server. Administrators want host-level backup capabilities with all the granular recoverability of traditional backup tools. They are looking for **single-pass backup** tools that can provide host, guest and application protection in as simple a way as possible, and, if possible, without the need to install software on every VM. However, administrators also want to be able to use the same technologies to protect both physical and virtual servers, and to be able to easily recover to similar or dissimilar virtual or physical environments.

A key requirement for any server virtualization protection strategy is to be able to provide fully granular recovery options. These would include the ability to recover individual files and folders, and specific application data such as Microsoft SharePoint objects, Microsoft Exchange emails or Microsoft SQL Server® tables, as well as complete VMs.

Irrespective of whether servers are physical or virtualized, in most organizations, email, file storage and database storage requirements have increased as the cost of disk storage has reduced. These large data volumes require large storage capacity, with much data being duplicated within or across storage volumes. For any backup or replication operation, particularly with the large amounts of data now common across virtualized server deployments, there is significant potential for performance hits on both the source and destination computer systems, as well as significant consumption of network bandwidth. It is, therefore, essential to be able to manage performance to minimize potential system performance degradation during data protection operations, and particularly when managing the disk I/O bottleneck.

In a virtualized server environment, it is important that backup and system protection, as well as recovery processes, are as automated as possible. Backup processes have a much more direct and immediate impact on virtualization infrastructure resources now than they did when each machine was distributed using separate physical hardware. The challenge now is managing the impact of parallel data and system protection on virtualization environments, using technologies such as automated backup and recovery, integrating snapshot-based disk backups with traditional long-term backups and integrating replication and high availability tools with backup. Backup and system protection technologies must also be able to support low-cost standby systems, with options for automated failover, if required.

# Backup and recovery

Backup and recovery refers to the protection of critical data and applications across your server environment, such as user files and folders on file servers, Exchange mailboxes, and SQL Server databases. Good virtual server backup and recovery must be able to provide options for backing up the whole system, or specific VMs or applications, and must be able to deliver consistent backups across all VMs on a host server in one pass. In the event of VM failure, or catastrophic failure at the host server level, recovery technologies must be able to provide rapid restores, even if the recovery is to a different virtual environment or to a host server using different physical hardware.

## CA ARCserve Backup r16

CA ARCserve® Backup creates and manages secure permanent data backups and provides bare-metal disaster recovery tools so that a computer without any operating system or application software can be restored to a previously backed-up state. Backed-up data can be stored on disks or tapes, or in cloud storage, and CA ARCserve Backup supports data migrations through staged backup jobs. For example, VMs can be regularly and rapidly backed up to local disk, and then this disk data can be copied or archived to tape or cloud for long-term, off-site storage.

CA ARCserve Backup supports two levels of backup for virtual servers:

1. By protecting the host server, all VMs and their virtual hard disks are automatically protected.
2. Using specific agents for each VM and application provides an additional level of protection, and enables more granular recovery of VMs and their applications.

CA ARCserve Backup provides agents for protecting physical and virtual servers running on Windows®, Linux® and UNIX®, as well as agents for popular applications. The same backup policies can, therefore, be used for both physical and virtual environments. CA ARCserve Backup agents installed in the host computer can provide both host-level and VM-level backups for Microsoft Hyper-V™ as well as VMware environments. For more granular recovery, agents can also be installed in the VMs of Hyper-V, VMware® and Citrix® XenServer platforms. For application-level protection, you can also install specific agents, such as the Agent for Microsoft Exchange Server, Agent for Microsoft SharePoint Server, Agent for Microsoft SQL Server and Agent for Oracle®, in order to get application-specific protection features such as mailbox and database recovery. Figure 1 shows how CA ARCserve Backup agents are deployed to protect physical and virtual resources.

**Figure 1. CA ARCserve Backup r16 backup sources and targets**

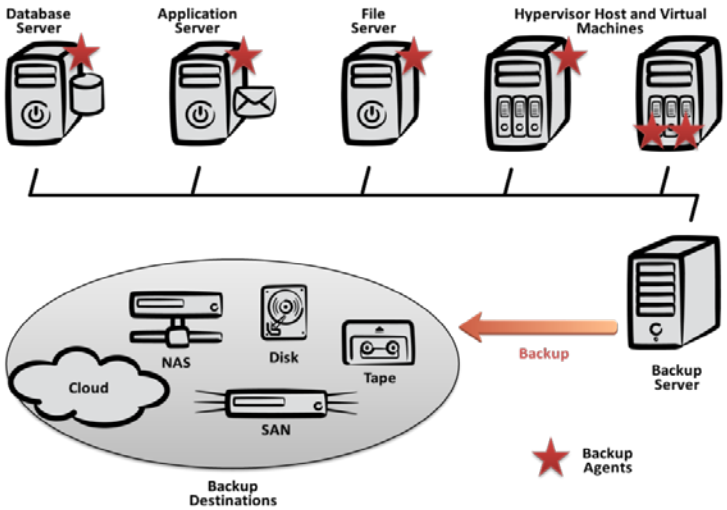

CA ARCserve Backup supports simultaneous backups of multiple VMs across Microsoft Hyper-V, VMware vSphere™ and VMware vCenter™ Server servers, and uses snapshot technologies to back up VMs, without affecting VM performance. CA ARCserve Backup also enables you to back up a VM whether the VM is powered on or off.

CA ARCserve Backup provides a range of options for granular restores. The CA ARCserve Backup Agent for Virtual Machines enables you to back up complete or raw VM images or file-level VM backups on supported Windows systems. Raw backups are more efficient when backing up an entire VM and can permit restore at a granular file level. Raw backups without file-level restore do not require agents to be installed on each VM. Installing the CA ARCserve Backup Agent for Virtual Machines on a VMware Consolidated Backup (VCB) or other proxy computer, or Hyper-V host computer, is all that is required. Using raw backups, you can restore full VMs, using the "Recover Virtual Machine" option. Alternatively, you can perform a granular file or folder restore by using the "Restore by Session" option, selecting the session number and browsing within the session to the VM files themselves, for example, *.vhd, *.bin and *.vsv files for Hyper-V VMs.

For more granular protection, installing the CA ARCserve Backup agent on each VM enables incremental and differential backups, including file and folder recovery from all backups, as well as raw backups. To provide the full range of granularity, you can enable mixed-mode backups, and include raw and file-level backups as part of the same job. Additional agents can also be installed at the VM level to provide application-aware backups for applications such as Microsoft SharePoint, Microsoft SQL Server and Oracle.

Data deduplication in CA ARCserve Backup reduces storage requirements by identifying data within volumes that is identical (Figure 2). The backup process ignores redundant data and stores only unique data to disk for subsequent backup. Using deduplication, you can fit more backups onto the same storage media, and retain backups for a longer time. To use deduplication, you back up your data to data deduplication devices (DDDs). These DDDs are used to store the actual unique data, as well as the reference and index files that enable CA ARCserve Backup to keep track of which files are made up of which unique data blocks. To maximize performance, the data and index files should be located on different disks from the actual data. When restoring deduplicated data from backup, CA ARCserve Backup uses the index files to identify and locate data segments before reassembling the original data stream.

**Figure 2. CA ARCserve Backup r16 deduplication in operation**

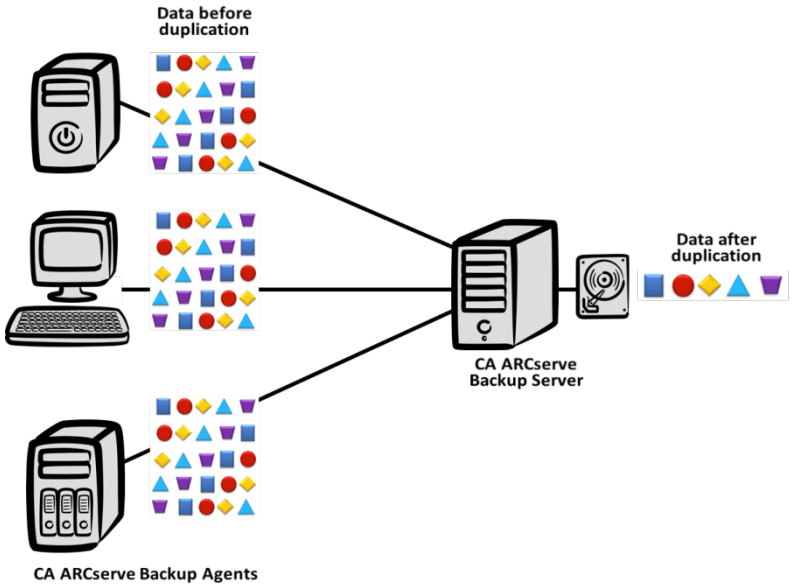


After CA ARCserve Backup agents have been deployed, the CA ARCserve Backup **Infrastructure Visualization** view provides a simple network diagram view of the entire physical and virtual environment, including all of the servers, storage and other devices. CA ARCserve Backup also provides auto-discovery and backup of VMs, helping to ensure that all VMs are backed up regardless of how the virtual environment develops and changes. Auto-discovery can be configured to run as frequently as every hour (the default is every 24 hours) in order to automatically discover new VMs. Live migration may be used, for example, to move VMs between hosts for performance reasons, or as part of planned downtime maintenance on host servers, but it is still important that the VMs continue to be backed up and protected.

CA ARCserve Backup includes integrated anti-virus, using the CA Anti-Virus r8.1 engine. This enables you to configure anti-virus scanning during backups, and provides an added layer of protection to make sure that you are not corrupting your backup with a virus.

Virus signatures can be updated automatically using the CA ARCserve Backup Job Scheduler Wizard, or manually from the command line.

CA ARCserve Backup also provides VM-specific tools to help avoid physical server bottlenecks when backing up VMs. For example, vSphere 4.x and Hyper-V andVSS integration technologies are used to offload backup processes and minimize the effect of backups on production VMs. For Microsoft Hyper-V, VMs show up in the CA ARCserve Backup Manager Console without any additional scripting or configuration. VMware VMs also integrate within the CA ARCserve Backup Manager Console through the VMware Virtual Disk Development Kit (VDDK) for VMware vSphere.

## How CA ARCserve Backup protects Hyper-V

CA ARCserve Backup supports VMs running on either stand-alone Microsoft Hyper-V Server 2008 and 2008 R2, or Windows Server® 2008 and 2008 R2 running the Hyper-V role. Full or Server Core installations are supported.

CA ARCserve Backup uses VSS Snapshots to enable point-in-time online backups of live VMs. CA ARCserve Backup integrates with the Hyper-V VSS writer so that backups are taken of VM snapshots and not the live VM itself while the data is consistent within the VMs themselves. The CA ARCserve Backup agent first initiates a VSS Snapshot of a running VM, and then backs up this snapshot, without any noticeable impact on the guest VM itself (Figure 3).

**Figure 3. Protecting Hyper-V using CA ARCserve Backup r16**

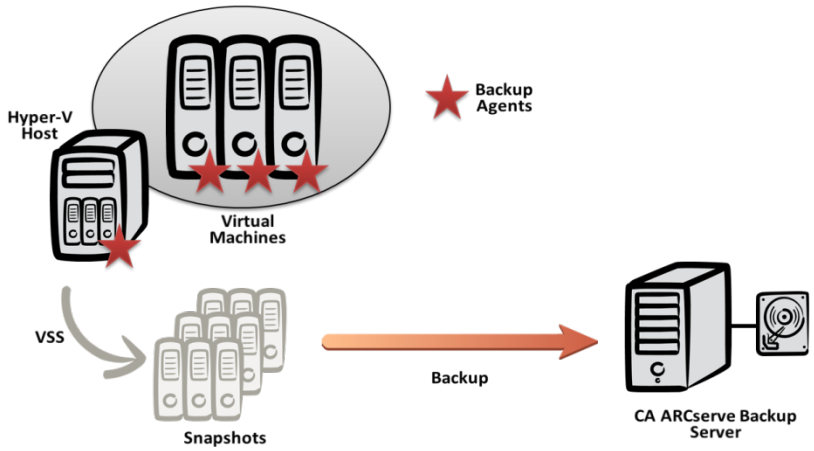


By using VSS, you can create a shadow copy that includes only the changes that have occurred since a full shadow copy was last completed, so the copy only takes up a small percentage of the overall volume size. By contrast, a hardware snapshot performs a full copy of a volume. This is an exact copy of the volume, so the copy requires the same amount of disk space as the volume. Hyper-V uses VSS writers to ensure that VMs are backed up to a consistent state when the shadow copy backup request is processed. When you create the shadow copy, the VSS writer flushes all data buffers and suspends writes to a volume to ensure that files selected for backup remain in a consistent state. For Windows virtualization, VSS support is provided through the Microsoft Hyper-V VSS writer. VSS writers are provided and supported by the application vendor, so any updates to the VSS writer are included with application updates. Therefore, by using this technology, CA ARCserve Backup ensures the consistency of backup data without requiring CA ARCserve Backup agents to be updated whenever new application versions are released. CA ARCserve Backup supports VSS through the Agent for Open Files and the Enterprise Option for VSS Hardware Snap-Shot. CA ARCserve Backup VSS support is automatically installed when you install the Agent for Open Files.

## How CA ARCserve Backup protects VMware

CA ARCserve Backup supports VMs running on VMware ESX Server 3.0 and above, and vSphere 4.0 and above, and uses a similar auto-discovery process as used in Microsoft Hyper-V environments.

For VMware running on ESX hosts in older VMware Virtual Infrastructure environments, CA ARCserve Backup enables you to offload VM backup activity to a dedicated backup proxy system, using VMware Consolidated Backup (VCB) integration features, and the proxy system can access and mount VM snapshots for backups (Figure 4). The backup proxy server takes periodic VM snapshots and it is these snapshots that CA ARCserve Backup then backs up from the backup proxy, with no impact on the ESX host.

**Figure 4. Protecting VMware using CA ARCserve Backup r16**

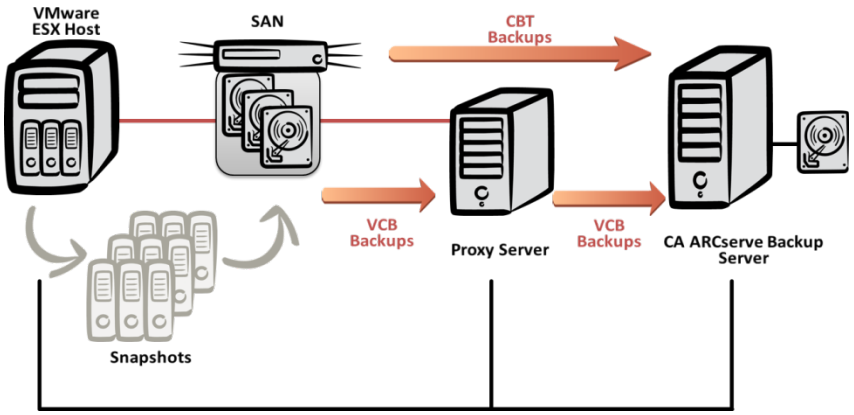


In VMware vSphere environments, CA ARCserve Backup uses vSphere integration features (provided through the VDDK), such as Changed Block Tracking (CBT), to provide fast snapshot backups without the need for backup proxies. CBT enables VMs running on VMware ESX 4.0 or later hosts (configured for virtual hardware version 7) to keep track of disk sectors that have changed. Information about these virtual disk block changes is tracked in the hypervisor (VMkernel), and CA ARCserve Backup accesses this information using VMware vStorage APIs for Data Protection (Figure 4). CA ARCserve Backup can make use of CBT for any type of virtual disk, thick or thin, and on any datastore type (NFS and iSCSI) except for physical mode Raw Device Mappings.

Using these VMware vSphere integration application programming interfaces (APIs), CA ARCserve Backup also enables other VMware-specific features such as scheduling VM backups to avoid contention with physical resources, direct data transfers from the VM to the backup device, and direct data restore from backups to VMs. The integration features in CA ARCserve Backup r16 eliminate the need for a staging server in VMware environments, and can also be used to back up data at the VMware storage level.

## CA ARCserve D2D r16

CA ARCserve® D2D provides disk-based protection for physical Windows servers and virtual Windows servers running on Hyper-V or VMware ESX hosts. Data and system information is protected as a recovery point and can be stored on local or network attached storage (NAS) disks and storage area networks (SANs), or files can be sent to off-premise storage. Recovery points can also be copied to cloud storage. CA ARCserve D2D provides single-snapshot backups of VMs with granular recovery capabilities; these backups enable quick restore of files, volumes, databases and emails as well as entire VMs. For example, if CA ARCserve D2D r16 Advanced Edition is used and the VM is running Microsoft Exchange, Granular Mailbox Recovery is automatically enabled. Similarly, if the VM is running Microsoft SQL Server, individual databases can be recovered.

To easily recover files and folders, CA ARCserve D2D provides Windows Explorer Shell Integration. When you right-click the folder containing the CA ARCserve D2D recovery points, and change to the “ARCserve D2D View,” Windows Explorer displays a list of all recovery points in that folder. Using this integration, you can easily open a recovery point, browse to the individual files and folders in that recovery point and then manually copy them for quick restores of individual files from a particular recovery point.

CA ARCserve D2D uses block-level Infinite Incremental (I2 technology™) snapshots for all backups. This technology automatically reduces storage requirements, requires fewer CPU resources during backups and reduces backup times. When you start a backup, the specified volume is divided into a number of subordinate data blocks that are then backed up. The initial backup is considered the "parent backup" and will be a full backup of the entire volume to establish the baseline blocks to be monitored. Prior to performing the backup, a VSS Snapshot is created. Then an internal monitoring driver checks each block to detect any changes. For all subsequent backups, CA ARCserve D2D will incrementally back up only those blocks that have changed since the previous backup. CA ARCserve D2D enables you to schedule the subsequent block-level incremental backups ("child backups") as frequently as every 15 minutes to always provide accurate, up-to-date backup images. You set a limit for the number of incremental child backups that are stored. When this specified limit is exceeded, the oldest incremental child backup is merged into the parent backup to create a new baseline image consisting of the "parent plus oldest child" blocks, synthesizing a full backup image by using the oldest incremental backup. (Unchanged blocks will remain the same.) This cycle of merging the oldest child backup into the parent backup repeats for each subsequent backup, enabling you to perform Infinite Incremental (I2) snapshot backups while maintaining the same number of stored (and monitored) backup images (Figure 5). In this way, there is no requirement for time-consuming post-backup consolidation of VM

backups, or for deleting older backup sets to make room for new ones. If you need to restore the volume information, the most recent backed-up version of each block is located and the entire volume is rebuilt using these current blocks.

**Figure 5. Infinite Incremental ($I^2$) backups**

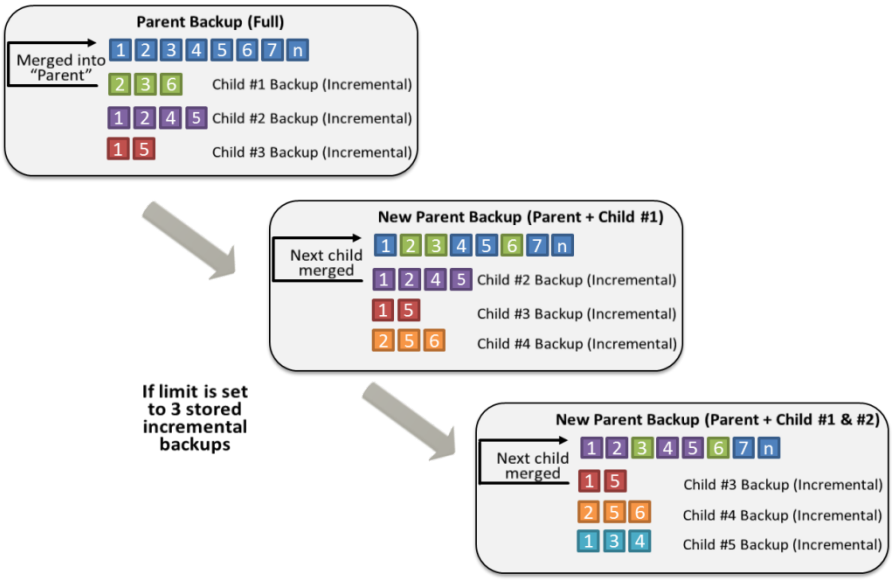


CA ARCserve D2D works with Windows operating systems that support the VSS writer. These operating systems are Windows Server 2003 Service Pack 1 (SP1) and later versions of the Windows Server operating system, Windows® XP and later versions of Windows desktop operating systems.

CA ARCserve D2D provides a **backup speed throttling** capability. You can use this to specify the maximum speed (MB/min) at which your backups are written to reduce CPU or network utilization. An important assessment exercise must therefore be to determine an appropriate level of throttling. As you increase the maximum backup speed, it reduces the amount of time to perform the backup. CA ARCserve D2D backup data can be compressed to save space, and encrypted using Advanced Encryption Standard (AES)-128, AES-192 or AES-256 for security.

You can also use **CA ARCserve D2D** backups for migrating Windows servers to, and from, virtual platforms using any of the following methods:

- **Bare Metal Recovery** (BMR) to similar or dissimilar hardware. This includes an option to resize destination volumes. Using BMR, you can recover a server using data that was backed up using CA ARCserve D2D.
- **Simple virtual conversion** (physical to virtual (P2V)) to enable backup of a physical server and restore to a virtual server.
- **Copy and scheduled export of recovery points** to take backed-up data and automatically export a copy to be taken off-site or to another location.

## How CA ARCserve D2D protects Hyper-V

By protecting the Hyper-V host, CA ARCserve D2D provides both host-level and VM-level protection. For Hyper-V host-level-only protection, you install CA ARCserve D2D on the Hyper-V host server. Then if the Hyper-V host server fails, you follow the standard BMR procedure in CA ARCserve D2D to recover the Hyper-V host server. If you want to restore selected files, you use the standard CA ARCserve D2D restore procedure and browse to the files you want to restore. Figure 6 shows where CA ARCserve D2D can be installed to protect Hyper-V VMs.

**Figure 6. Protecting Microsoft Hyper-V using CA ARCserve D2D r16**

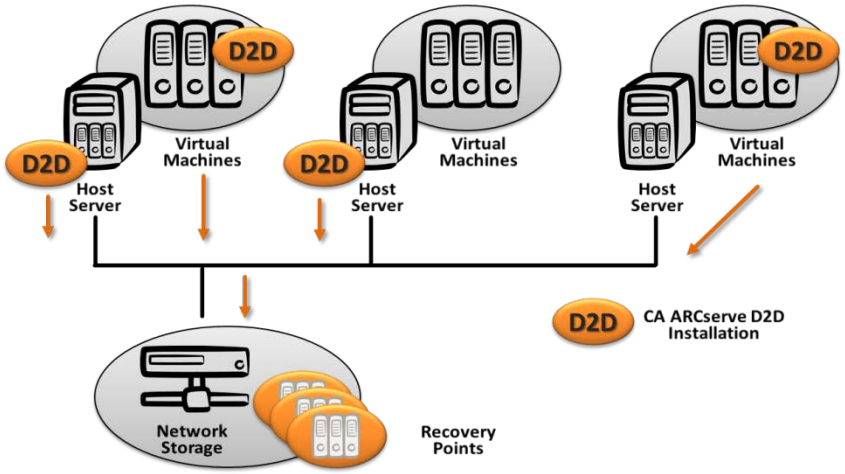


For Hyper-V host-level and VM-level protection, you install CA ARCserve D2D on the Hyper-V host server. Then to restore individual VMs from the CA ARCserve D2D backup, you select to restore the VM to the original location or an alternative, and then select the individual VM files (*.vhd, *.avhd, *.vsv and *.xml configuration files) from the restore window in CA ARCserve D2D. You can also install CA ARCserve D2D inside individual Windows VMs if you want to protect VMs separately from the host. This method is also needed if you attach an iSCSI logical unit number (LUN) directly inside the VM, because data inside the LUN will not be backed up using CA ARCserve D2D Hyper-V host-level backups.

### How CA ARCserve D2D protects VMware

In a VMware environment, CA ARCserve D2D can be installed on individual VMs, and used to protect VMs using the same approach as for any physical server. CA ARCserve D2D can also be used together with **CA ARCserve® Central Host-Based VM Backup**. This technology provides the ability to perform a single-pass backup of each VM on a VMware host server without having to install any agent or other software inside each guest VM. CA ARCserve Central Host-Based VM Backup protects VMs running on VMware vSphere 4.0 or later systems, and can recover data at VM level, application level and file level from a single-pass VM backup (Figure 7).

**Figure 7. Protecting VMware using CA ARCserve D2D r16**

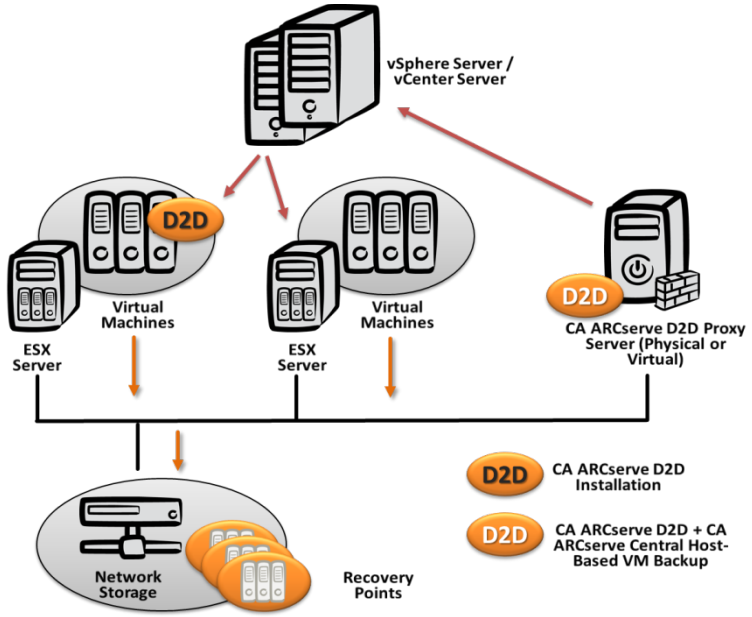


CA ARCserve Central Host-Based VM Backup uses the VMware vSphere CBT feature to capture the changes made to the VM since the last backup and enables granular recovery of files and folders from each individual VM. It can also restore supported applications running within a VM, simplifying protection and recovery of virtual servers. Backup sessions are stored in CA ARCserve D2D recovery point format, and can be accessed by CA ARCserve D2D and by the CA ARCserve D2D Bare Metal Recovery (BMR) tools to restore VMware ESX VMs either to other VMs or to a physical server.

CA ARCserve Central Host-Based VM Backup automatically discovers recently added virtual machines, as long as those VMs are running a Windows operating system. CA ARCserve Central Host-Based VM Backup requires CA ARCserve D2D to be installed on the backup proxy system.

### How to manage CA ARCserve D2D deployments

CA ARCserve D2D uses a Web 2.0 interface for configuring and managing backup and recovery tasks. CA ARCserve D2D can be remotely deployed over the network using the CA ARCserve D2D console, with nodes added manually by name. Once it is deployed, you can select these remote nodes from the base CA ARCserve D2D homepage for management. This interface can be used to manage individual CA ARCserve D2D servers, but additional tools from the CA ARCserve® Central Applications suite are particularly useful where there are large CA ARCserve D2D deployments:

- **CA ARCserve® Central Protection Manager** provides an alternative to the CA ARCserve D2D console, with easy access to all CA ARCserve D2D backups across the network. It can be used to restore files, folders and applications from all CA ARCserve D2D recovery points. CA ARCserve Central Protection Manager also adds the ability to auto-discover physical and virtual CA ARCserve D2D servers by using computer objects stored in Active Directory®.
- **CA ARCserve® Central Reporting** is used to collect information and view reports about the performance of CA ARCserve D2D nodes and CA ARCserve Backup servers from a central location. You can view reports in tabular and chart formats in a browser-based, dashboard interface, and filter data to view reports about specific branches or groups of protected computers so that you can target report data that is unique to a set of systems with common characteristics. For example, the ***Virtualization Protection Status Report*** displays the status of all VMs backed up using VCB or Hyper-V. Data can be exported as CSV files or sent via email.

### How CA ARCserve D2D works with CA ARCserve Backup

Using CA ARCserve Backup, you can import CA ARCserve D2D recovery points from multiple CA ARCserve D2D servers and store the data on CA ARCserve Backup media such as tapes. To restore from a single CA ARCserve D2D VM backup using the CA ARCserve Backup Manager Console, you use the "Restore by Tree" option and select the CA ARCserve D2D Server RAW Session. After you have restored the whole CA ARCserve D2D session to an alternative location, you use the restored sessions to recover the CA ARCserve D2D data. Using CA ARCserve Backup, you can recover CA ARCserve D2D data at file, folder and application level, including Microsoft SQL Server and Microsoft Exchange Server application data. Using CA ARCserve D2D with CA ARCserve Backup helps reduce or eliminate backup windows. Using I2 technology, CA ARCserve D2D recovery points are quick to create. CA ARCserve Backup can then manage and move these recovery points offline with no impact on the backed-up VM or host server.

You can also use CA ARCserve Backup backups of CA ARCserve D2D data to perform Bare Metal Recovery (BMR) of CA ARCserve D2D servers. Using the raw backup sessions, the BMR process is a two-phased approach:

1. Recover the raw session to a shared folder, a network file share or a device that the failed server can access while completing the BMR process.
2. Boot the server that you want to recover using the CA ARCserve D2D BMR media and then browse to the location where you recovered the raw session. Follow the on-screen steps to complete the BMR process.

# Replication

Replication refers to the real-time continuous protection of complete systems and data, as well as the ability to copy or migrate backups and data to one or more servers and storage at remote sites for disaster recovery purposes.

## CA ARCserve Replication r16

CA ARCserve® Replication complements CA ARCserve Backup and CA ARCserve D2D, and almost any other backup solution, by providing enhanced data protection through continuous data replication, off-site data protection for disaster recovery (DR) and Data Rewind for CDP. Using CA ARCserve Replication, file and database changes are automatically copied in real time from a Production (live) server to a physical or virtual Replica server. Replication also eliminates VM storage as a potential single point of failure by ensuring that there is a copy of all VMs on alternative storage, which can include local, off-site or cloud storage. Even if you already use technologies such as Windows failover clustering, CA ARCserve Replication provides valuable additional protection through off-site replication. CA ARCserve Replication protects VMware ESX, vSphere and Citrix XenServer at the VM level, and protects Microsoft Hyper-V at the hypervisor and VM levels. It can also protect Windows and Linux systems on physical and virtual servers.

The CA ARCserve Replication management module, with its control service and associated user interface, enables you to easily deploy replication modules or engines to the Production and Replica servers. CA ARCserve Replication includes central deployment, management, reporting and maintenance tools. It is not necessary to manually install the CA ARCserve Replication engine on the Production or Replica servers. This is because, as part of the Replication scenario deployment, the engine can be automatically deployed and this installation does not require the Production or Replica server to be rebooted.

After the CA ARCserve Replication control service and engine has been installed, you create scenarios to protect your VMs and applications. For example, to protect a Hyper-V host server and all its VMs, you create a **Microsoft Hyper-V Replication and Data Recovery** scenario. After selecting the Production and Replica host servers to use with the scenario, you can choose to verify the CA ARCserve Replication engine on the host servers (and deploy if the engine has not been installed or requires updating). All VMs running on the Hyper-V Production server are then auto-discovered and, by default, all VMs are selected for replication.

For application-level protection, CA ARCserve Replication uses auto-discovery of the application environment to ease and speed deployment. You can use auto-configuration to create Replication scenarios to protect Microsoft Exchange, Microsoft SQL Server, Microsoft SharePoint, IIS, Microsoft Dynamics® CRM, and Oracle. You can protect other Windows applications by using the Custom Application Protection Wizard or custom scripts. Real-time continuous data protection (CDP) provides data replication to any local or remote physical or virtual server and storage for disaster recovery purposes.

For all scenario types, after an initial synchronization between the Production server and the Replica server has completed, the CA ARCserve Replication engine only sends byte-level changes to the Replica server. This technique reduces the bandwidth that is required for daily backups of remote data and applications. For protecting VMs, block-level synchronization is used. The CA ARCserve Replication engine performs a block-by-block comparison of the files on the Production and Replica servers, and copies over only those blocks that are different. In this way, when differences exist between large VHD files, instead of requiring the transfer of the entire VHD file, block synchronization transfers the changes only. To speed up deployment, you can use **Offline Synchronization** to import data to the Replica server from physical media. This can be especially useful for initial replication over the wide area network (WAN).

CA ARCserve Replication also provides additional performance-related technologies. **Multi-stream replication** means that replication data can be sent over multiple IP sessions even within a single scenario. This improves replication and synchronization times for most scenarios, but has the greatest impact for customers with scenarios running across high-latency WAN connections. **Bandwidth throttling** enables you to control the size of the incoming bandwidth allowed on the Replica host. You can either define one size limit that will apply to all hours of the day, or you can specify different values for different hours. By using the bandwidth scheduler, you can decrease the bandwidth size during busy hours and increase it during off-peak hours in order to optimize your bandwidth resources.

CA ARCserve Replication also provides an **assessment mode** that enables you to measure the amount of bandwidth that is required to replicate data, prior to implementation. Assessment mode enables you to plan your bandwidth requirements and adjust either your bandwidth or the volume of data that you want to replicate to suit your individual requirements.

CA ARCserve Replication provides automatic recovery testing called **CA ARCserve® Assured Recovery®** to ensure that your replicated data and applications are recoverable. After the CA ARCserve Assured Recovery test completes, you can use the Microsoft Volume Shadow Copy Service to generate a snapshot of your data and application to provide additional protection for your critical data and enable a restore point. VSS Snapshots enable you to create a point-in-time image copy of data on a volume. By mounting a particular snapshot as a file system device using the CA ARCserve Replication console, you can quickly restore individual files or entire volumes in the event of system failure or data corruption.

CA ARCserve Replication supports the scheduling of VSS Snapshots on the Replica server. This offers another failback option to recover data. These snapshots can then be mounted for recovering whole volumes, or individual files and folders, to a specific point in time. CA ARCserve Replication provides a Custom Application Protection Wizard for Windows-based applications to help eliminate the need for scripting an application-specific failover or recovery processes, such as to protect applications that do not have their own built-in scenario.

CA ARCserve Replication includes Data Rewind for CDP, to provide fast recovery of lost or damaged data and databases to a specific point in time. Data Rewind uses rewind journals to store the I/O operation information that results in modified files. Using the rewind journal, it is possible to undo I/O operations, and rewind the file to a previous point in time, such as to a valid, non-corrupted state. Data Rewind complements a backup solution and protects data between periodic backups for faster recovery time and more granular recovery points, but does require the Replica server to be online.

CA ARCserve Replication also provides a **Full System** replication scenario. Using this scenario, an entire system is replicated including the operating system, system state, application and data from any physical or virtual server to an offline virtual server that supports the guest operating system of the active server. Full System replication is capable of replicating a physical or virtual machine to three different virtual server formats: Hyper-V, VMware ESX and Citrix XenServer. Replication can also be to Amazon Elastic Compute Cloud (Amazon EC2), for Windows-based virtual machines hosted in the cloud. By saving replicas to offline storage, you may not require the additional licensing of the VM operating system and applications as would be required for hot standby servers that were running live. It also enables potential over-provisioning of your virtual server destinations, which are only used when required for disaster recovery. The offline Replica server can be made available if the Production server is down. The Replica VM is effectively a clone of the Production server, so it must remain in an offline state until failover to prevent IP, name and/or network conflicts.

CA ARCserve Replication also enables P2V and virtual-to-virtual (V2V) migration. You can quickly and easily migrate from physical to virtual servers, and migrate between virtual server platforms (such as VMware to Hyper-V) as part of a virtualization consolidation project, for example.

#### How CA ARCserve Replication protects Hyper-V

For Hyper-V, CA ARCserve Replication provides flexible deployment options. You can choose to replicate at the hypervisor level or to replicate individual guest operating systems and applications without the requirement of installing the CA ARCserve Replication engine on each of the guest VMs.

In order to protect a Hyper-V environment, you can create one Hyper-V scenario per Hyper-V Server. CA ARCserve Replication will automatically generate an individual scenario for each Hyper-V guest on the servers. For example, if there are 10 Hyper-V VMs on Server A, you would use the Scenario Creation Wizard to create a scenario for Server A to replicate to a standby Hyper-V server. CA ARCserve Replication will automatically generate a Protection scenario for each Hyper-V VM on Server A to facilitate failover of individual VMs. So, in this example, 10 scenarios would be created for the VMs running on Server A, and these scenarios would be part of the same Scenario Group called “Server A”; all of this is done automatically. When using CA ARCserve Replication at the hypervisor level, some features may be unavailable or have limited functionality. For example, the CA ARCserve Assured Recovery feature is unavailable and the Data Rewind function offers reduced granularity when selecting a rewind point. When these features are required, a hybrid scenario can offer the best of both hypervisor and guest-level replication. A hybrid scenario is one where you might identify a handful of VMs or guests, such as a server running Microsoft Exchange or SQL Server, that would benefit from having CA ARCserve Assured Recovery and the more granular Data Rewind features available. These VMs would then get their own Replication scenarios and CA ARCserve Replication engine. All other VMs would be configured as part of the hypervisor-level Protection scenario. Alternatively, you can create scenarios at an individual VM level.

#### How CA ARCserve Replication protects VMware

To protect VMware VMs using CA ARCserve Replication, you need to create separate file server, application or Full System scenarios for each VM, and deploy the CA ARCserve Replication engine on each VM. There is no host-level replication option.

In a VMware environment, you should make sure that your VMware vCenter infrastructure is protected in order to ensure that you can continue to manage all your VMs. Using the **VMware vCenter Server** scenario, you can use CA ARCserve Replication to replicate all the VMware vCenter components (Database Server, License Server and Web Access Server), whether they are installed on a single VMware vCenter server or are distributed across multiple servers.

# High availability

High Availability refers to the real-time protection of complete systems or applications, and where a backup system can be instantly brought online either manually or automatically.

## CA ARCserve High Availability r16

CA ARCserve® High Availability provides the same features as CA ARCserve Replication, and adds both system-level and application-level monitoring, automatic and push-button failover and push-button failback.

To deploy this solution, you can build your own physical or virtual Replica servers and then synchronize and replicate the data to those Replica servers that may be located on site or at any remote location. Alternatively, you can replicate complete physical or virtual

systems, including the operating system, system state, application and data, using Full System Protection, to an offline virtual Replica server that may be located on site or at any remote location.

Failover and failback describes the CA ARCserve High Availability process in which active and passive roles change between Production and Replica servers. Failover is used if there is a problem with the Production server, and can be automatic or manual ("push-button"). Failback is used after the Production server has been repaired or replaced, in order to get the original Production server resynchronized with the current Production server (the original Replica server before the failover). Failover and failback support application failover on 32-bit and 64-bit systems running on the VM, just as if they are physical machines. You can configure failover between Production and Replica servers to occur automatically when the CA ARCserve High Availability control service detects that the Production server or application is unavailable, or manually through the CA ARCserve High Availability Manager. You can configure automatic failover on VMs by using predefined monitoring checks, which include ping checks, database checks and user-defined checks. User-defined checks enable you to tailor the failover to meet specific application requirements.

Push-button failovers, with automated network and end-user redirection, can be used in advance of an impending disaster or outage. They are also a good solution for "hot" migrations that can be performed live during the working day with minimal disruption to users. At the end of the migration, the new server is immediately ready to go live.

CA ARCserve High Availability also includes CA ARCserve Assured Recovery to provide automated, non-disruptive "lights-out" recovery testing that can be scheduled for recurring testing.

CA ARCserve High Availability also provides a "Failover to Cloud" (Full System) scenario, which is specific to Amazon EC2, for Windows VMs, where the Production server is a local physical or virtual server and the Replica server is an Amazon EC2 server. In order to provide seamless failover using the cloud, you can specify to automatically **redirect DNS** with this scenario, so that user requests for the VM are automatically redirected to the Amazon EC2 server using your own Amazon Virtual Private Cloud (Amazon VPC) configuration. The Amazon VPC is a private, isolated section of the Amazon Web Services™ (AWS) cloud where you define your own virtual network topology, usually in such a way that it is similar to your own data center topology. Please refer to the technical white paper called "Leveraging the Cloud for System and Data Backup, Recovery, Availability and Archiving with CA ARCserve® r16 Products" for more information.

CA ARCserve High Availability also enables you to create distributed groups. This is particularly useful for SharePoint server farms or other application environments where the integrity of a service depends on multiple physical or virtual servers. Unlike the default group and the regular group, the distributed group has central management functionalities. Each server still needs its own scenario, but there are now common scenario properties that apply to the whole group and do not need to be separately specified for each scenario. Some of the central management functionalities include:

- **Group Run and Group Stop**. You can start and stop all the scenarios in the group together.
- **Group Failover**. You can initiate manual failover on all the scenarios once, and configure them to be switched over together automatically in case any of them is failed.
- **Group Recover Active Server**. You can resolve the split-farm problem (some SharePoint server Masters are active while other Replica servers are active). It can easily recover active servers of all scenarios to either Production or Replica servers.

### How CA ARCserve High Availability protects Hyper-V

To protect virtual servers, you must install the CA ARCserve High Availability engine onto a Production Hyper-V server and on the standby (Replica) Hyper-V server in your environment. By installing the CA ARCserve High Availability engine on a single Hyper-V server, your protection is limited to only enabling data replication of the Hyper-V VMs to a non-Hyper-V server. To enable failover at the Hyper-V host level, the CA ARCserve High Availability engine must be installed on a second Hyper-V server. To enable failover at the VM level, you must also install Hyper-V integration components in each guest operating system. Failover can be configured to be automatic or manual. Where integrity checking is fundamental prior to failover, CA ARCserve High Availability enables you to failover virtual servers manually. The Hyper-V scenarios in CA ARCserve High Availability protect the entire VM and failover will occur at the VM level (Figure 8). If you want to failover at the application level (SQL Server, Exchange, Oracle and so on), you must install the CA ARCserve High Availability engine on the Production and Replica Hyper-V VMs and create the appropriate scenario type.

**Figure 8. Hyper-V failover using CA ARCserve High Availability r16**

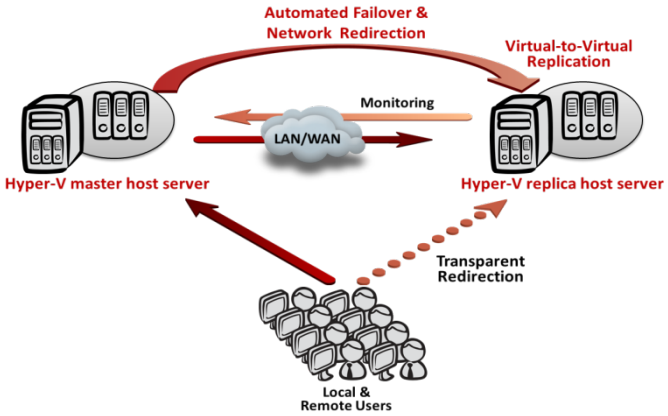


## How CA ARCserve High Availability protects VMware

For VMware VMs, CA ARCserve High Availability includes replication and high-availability support for VMware vCenter Server installations managing one or more VMware ESX servers, without the need for a clustering architecture or a shared storage infrastructure. To protect your VMware vCenter infrastructure, you must install the CA ARCserve High Availability engine on each VMware vCenter server in your environment (Figure 9). In a distributed environment, where the VMware vCenter databases are stored on separate machines, the CA ARCserve High Availability Production and Replica servers must point to the same distributed database. You should then protect these remote database servers by using CA ARCserve High Availability for SQL Server or CA ARCserve High Availability for Oracle.

**Figure 9. VMware vSphere and VMware vCenter failover using CA ARCserve High Availability r16**

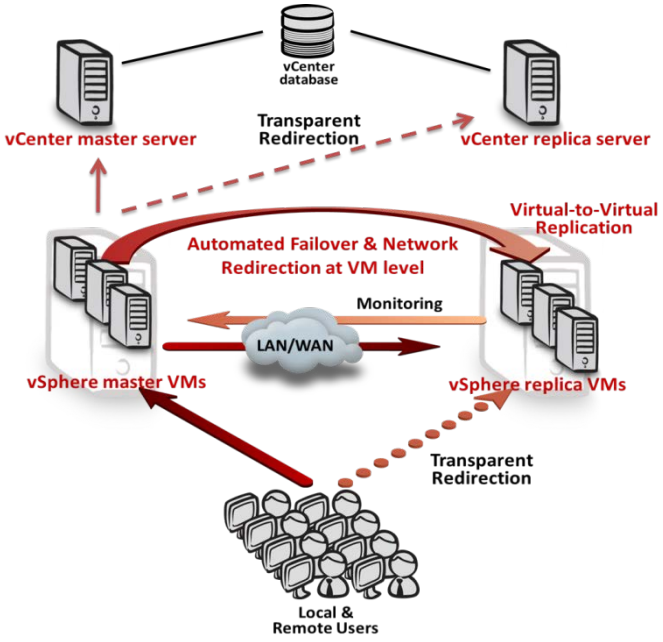


To replicate the actual VMware VMs, you use separate file server or application scenarios for each VM, and install the CA ARCserve High Availability engine on each VM. As with Hyper-V, you can configure failover between the Production and Replica servers to occur automatically when the Production server is unavailable, or manually through the CA ARCserve High Availability Manager. In a VMware environment, this failover configuration applies to both the VMware vCenter Protection scenario, and to the individual VM scenarios. By contrast, VMware vCenter Site Recovery Manager uses image cloning when provisioning a VM from one machine to another, so that the new VM will have the same machine name and IP address as the failed VM. Using this approach, it is very difficult to remotely access the Production VM to understand why it failed, troubleshoot the system and restart the failed application. In CA ARCserve High Availability, failover occurs by redirecting the server traffic and users at the DNS server level. This means that no changes are made to any server configuration of the Production or Replica VMs during or after a failover. This also makes managing VMs off site much

easier because it is not necessary to create a virtual local area network (LAN) to enable an IP address from your Production subnet to work on the DR site (which is most likely to be on a different IP subnet).

## CA ARCserve Central Virtual Standby r16

In some cases, a true high-availability solution may not be warranted, but backup and restore, or even Bare Metal Recovery (BMR), may take too long to meet recovery time objectives (RTOs). For these situations, using CA ARCserve D2D with CA ARCserve® Central Virtual Standby may be the right solution. CA ARCserve Central Virtual Standby works with CA ARCserve D2D to enable standby or backup Hyper-V or VMware VMs to be powered on manually or automatically using the VMDK or VHD files if the Production server is unavailable, or if you simply want to put a master or source server offline for maintenance. CA ARCserve Central Virtual Standby uses CA ARCserve D2D snapshots, or recovery points, and automatically converts these recovery points into VM snapshots based in configurable conversion policies and schedules. For fully automated standby, a heartbeat is used to detect the availability of the Production server. If the CA ARCserve D2D monitor server (this is the hypervisor host for Hyper-V) fails to communicate with the Production server within the heartbeat interval period, the host automatically starts the standby copy of the VM. This technology provides "ready-to-go" VMs for recovery without the disk or network overhead of restoring data from disk, tape or cloud-based backups. You can also recover data from the recovery point snapshots to the original or alternative source servers to perform virtual-to-physical (V2P) recoveries.

In addition, you can use the snapshots to perform V2P disaster recovery from virtual machines to the original or different hardware to recover to another VM (V2V) on another host (this does not need to be running the same hypervisor). You can reduce risk during physical-to-virtual (P2V) migration operations by working from a backed-up copy of the physical server state.

CA ARCserve Central Virtual Standby is not primarily designed to protect VMs, but to provide standby VMs as a recovery option for physical servers or for Production servers running as VMs. However, CA ARCserve Central Virtual Standby backups can be recovered using CA ARCserve D2D tools and individual files and folders can be accessed if required.

# Using the CA ARCserve product family to protect virtual environments

Virtualization vendors, such as Microsoft and VMware, provide some data protection tools that can be used in virtual environments, including Microsoft System Center Data Protection Manager, Microsoft Exchange database availability groups (DAGs), Windows failover clusters, VMware clusters and VMware SAN replication. CA Technologies, however, is independent of the virtualization technology and provides integrated data and system protection technologies that work across physical and virtual environments, across Microsoft servers and applications and on Microsoft Hyper-V, VMware vSphere, VMware ESX and Citrix XenServer hypervisors. In addition, CA ARCserve products can be used on their own, with other vendor products or as a complete solution to help meet many recovery point and recovery time objectives by application, server or data set.

CA ARCserve products support a range of virtualization-related integration scenarios such as:

- You can use **CA ARCserve Replication** to replicate **CA ARCserve D2D** (and **CA ARCserve Backup**) backup images to a remote location and the cloud, or to other media, for added protection. For example:
  - Use **CA ARCserve D2D** to back up all remote office or branch office VMs to local disk, and then use **CA ARCserve Replication** to replicate the backup files to a central location for off-site storage.
  - Use **CA ARCserve D2D** to back up VMs and files, and then use **CA ARCserve Backup** to migrate the CA ARCserve D2D backups to tape for long-term off-site retention.
- You can use **CA ARCserve Replication** to replicate all data to a central location and then use **CA ARCserve Backup** or **CA ARCserve D2D** for centralized backup. By performing VM backups from the Replica server, you can eliminate backup window constraints and performance issues. Using this integration and adding a CA ARCserve Backup property to Replication and High Availability scenarios to enable combined backup and replication jobs, you get CDP from CA ARCserve Replication so that data should never be lost, even if the host server fails. You also get long-term permanent storage to meet archival and compliance requirements. Virtual servers can be used as the Replica servers to reduce costs.

- You can use **CA ARCserve High Availability** combined with **CA ARCserve Backup** or **CA ARCserve D2D** to get high availability and total continued data protection, and to help overcome backup window constraints. Using CA ARCserve High Availability and CA ARCserve Backup together also enables you to restore data without an impact on your business operations. CA ARCserve Backup restores data to CA ARCserve Replica servers without affecting live Production servers at all. The Replica server and live (master) server are then resynchronized to complete the restoration process. Restoring to a Replica server also enables you to test your ability to restore data to your live environment.
- You can use **CA ARCserve Backup** to list backup-enabled **CA ARCserve Replication** and **CA ARCserve High Availability** scenarios as backup sources. You can then back up these sources by using the same simple steps as for standard backups. These scenario backups also automatically create VSS Snapshots of data on the Replica server, which you can save for permanent storage.

# Summary

Few organizations can standardize on a single server deployment environment, or single virtualization or guest operating system vendor. In most cases, there is a mix of physical and virtual servers, Microsoft Hyper-V, VMware and other hypervisors and Windows and other operating systems. It is important that data and system protection strategies can effectively cope with mixed environments, and help eliminate the potential single points of failure in a virtualized server infrastructure.

CA ARCserve technologies can be used together, or individually, to provide complete protection for virtual and physical servers, for backup and recovery of applications and data, for disaster recovery and replication of complete systems or applications and for ensuring high availability of mission-critical services.

For more information on the CA ARCserve Family of Products, please visit arcserve.com.