

## EMC Avamar

### The Big Picture

- Optimizes data deduplication at the source before transfer across networks
- Reduces daily full backup times by up to 10x
- Reduces required daily network bandwidth by up to 500x
- Reduces total backup storage by up to 50x
- Encrypted, electronic daily backups via existing LAN/WAN networks
- High availability and reliability; data recoverability and server integrity verified daily
- True centralized, multi-site backup control via intuitive web-based management
- Immediate, simple one-step recovery; no need to restore last full plus incremental backups
- VMware infrastructure backups: Guest, ESX Service Console, VMware Consolidated Backups (VCB)
- Flexible deployment options: EMC Avamar Data Store, EMC Avamar software only, EMC Avamar Virtual Edition for VMware
- Ideal for VMware /virtualized environments, remote and branch offices, and problem areas in data centers

## Next-generation backup and recovery with global, source data deduplication

Exponential data growth, regulatory compliance, strict service-level agreements, and shrinking backup windows are forcing enterprise organizations to rethink their data protection methods. And with companies pursuing aggressive virtualization strategies and struggling to protect data residing at remote offices, IT managers face a new set of challenges.

EMC® Avamar® is backup and recovery software with integrated global, source data deduplication technology. Avamar solves the challenges associated with traditional backup by delivering fast, daily full backups for VMware® environments, remote offices, and LAN/NAS servers in the data center. Unlike traditional backup methods, Avamar eliminates redundant sub-file data segments at the source (client) before data is transferred across the network and stored to disk. As a result, the required daily network bandwidth is reduced by up to 500x, enabling fast daily full backups across existing WAN/LAN links and virtual infrastructure. Avamar also deduplicates backup data globally, across sites and servers, reducing required total back-end disk storage by up to 50x. Backup data can be encrypted in flight and at rest, enabling secure, cost-effective, long-term retention on disk.

### Traditional data protection challenges

One of the key drivers impacting backup performance is the amount of data that must be protected within the available backup window. Traditional solutions are inefficient because they back up everything—including duplicate data files and sub-file data segments that exist across thousands of servers, desktops, laptops, and offices worldwide. When combined with traditional daily incremental and weekly full backup schedules, the impact of duplicate data is staggering—companies often move more than 200 percent of their primary data every week. As a result, meeting short backup windows can be very difficult due to the sheer volume of data that must cross already congested networks, backup servers, and infrastructure.

The impact is especially severe for virtual environments, remote offices, and NAS filers. In virtual environments, each virtual machine represents an individual backup job, often with concurrent or overlapping backup windows, and includes redundant operating system, application, and file data. As a result, backups for virtual machines can often overrun backup windows and tax shared resources, leaving data unprotected and creating management issues for backup administrators. In remote offices, limited network bandwidth makes centralized, automated WAN-based backup nearly impossible. Instead, firms rely upon non-technical local staff, failure-prone tape-based hardware, and ad-hoc manual processes that often result in unreliable remote office data protection. Protecting NAS filers can be a significant challenge too, especially when full backups fail to complete within the allotted timeframe, which can hurt employee productivity and leave data unprotected.

By backing up duplicate data, traditional solutions also increase costs due to the extra storage capacity required. This is exacerbated by the need to retain data for many months or years in support of archive or regulatory compliance objectives. In addition, traditional backup often involves the shipment of physical tapes offsite, including unencrypted tapes, which can result in exposure of confidential information, theft, or data loss.

**“Tape just wasn’t doing the job for us. The Avamar solution gives us more-effective data protection and puts us in a better position to restore lost data if necessary. In turn, we’re able to recover failed systems faster and shorten time to recovery for customer-facing services.”**

Todd Gourd  
Manager of IT Systems Administration  
Cherokee Nation Enterprises

## Game-changing next-generation backup and recovery

EMC Avamar enables fast, efficient backup and recovery by reducing the size of backup data at the source—before it is transferred across the network and stored to disk. Unlike traditional solutions, Avamar leverages the existing physical and virtual infrastructure to deliver fast, daily full backups. Avamar also deduplicates backup data globally across servers, desktops, notebooks, and offices worldwide to reduce the total required disk storage by up to 50x. As a result, Avamar enables efficient long-term retention of backup data on disk while dramatically lowering capital and operating expenses including floor space, power, and cooling.

Unlike traditional solutions, Avamar provides daily full backups that can be quickly recovered in just one step—eliminating the hassle of restoring full and subsequent incremental backups to reach the desired recovery point. And backup data is encrypted during transit across the network and at rest for security.

In addition, Avamar’s centralized web-based management and at-a-glance dashboard view makes it easy and efficient to protect hundreds of offices worldwide from a single location over existing network bandwidth.

## Avamar: efficient data deduplication

The method for determining segment size is a key factor in eliminating redundant data at a sub-file level. Some solutions on the market utilize fixed-block or fixed-length segments when performing data deduplication. With this approach, even small changes to a dataset (for example, inserting data into the beginning of a file) can change all subsequent fixed-length segments in a dataset. So, despite the fact that very little of the data has actually changed, the entire file will appear as new data that must be backed up again. Avamar solves this problem by examining the data to determine logical boundary points using variable-length data segments. As a result, Avamar delivers the most-efficient global, source data deduplication on the market—dramatically reducing the amount of data sent and stored, while increasing performance.

## Scalability, high availability, and reliability

Unlike many physical server deployments, the Avamar server utilizes a grid architecture that enables linear performance increases by simply adding storage nodes. Each incremental node increases CPU, memory, I/O, and disk capacity for the entire grid. When additional storage is added, data is automatically load-balanced without compromising deduplication efficiency or system performance.

When traditional backup solutions fail, enterprises are exposed to windows of potential data loss. Avamar employs patented redundant array of independent nodes (RAIN) technology to provide high availability and fault tolerance across nodes in an Avamar grid, eliminating single points of failure. In addition, Avamar system integrity is verified twice daily via internal system checkpoints. And Avamar verifies the recoverability of all backup data daily—so there are no surprises.

## Flexible deployment models to fit your exact needs

Avamar offers flexibility in solution deployments, depending on the specific use case and recovery requirements.

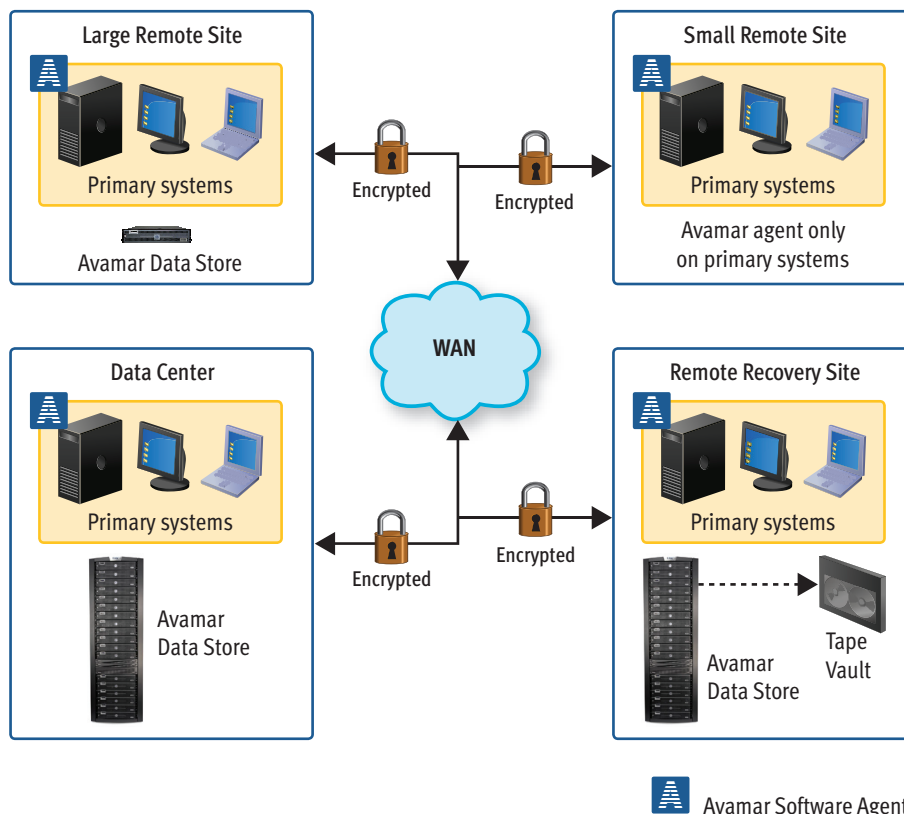
There are two convenient physical deployment options—EMC Avamar Data Store, a complete, pre-packaged backup and recovery solution that integrates Avamar software with EMC-certified hardware for streamlined deployment and EMC Avamar software that can be deployed on a range of certified industry-standard servers. An entry-level Avamar Data Store configuration is ideal for remote offices when fast, local backup and recovery are priorities. For larger remote sites and data centers, a range of Avamar Data Store configurations provides the ability to retain the equivalent of petabytes of daily full backups that can be immediately recovered in a single step.

For environments that have standardized on VMware virtual infrastructure, EMC offers a virtualized deployment option. The EMC Avamar Virtual Edition for VMware, the industry’s first deduplication virtual appliance for backup and recovery, consists of EMC Avamar software deployed as a virtual appliance. It enables a complete Avamar server to be easily deployed on an existing ESX Server™, leveraging the attached disk storage and infrastructure for a fast return on investment. Since all aspects of the backup and recovery process are encapsulated and virtualized, control and management

EMC Avamar Data Store



## EMC Avamar Flexible Deployment



are also streamlined, reducing demands on IT staff. Avamar Virtual Edition has the added benefit of cost-effective Avamar virtual-to-virtual, or Avamar virtual-to-physical server replication to meet disaster recovery objectives.

**“[EMC] Avamar provides a huge time and cost savings for backup of multiple virtual machines.”**

Curtis Damhof  
Senior Network Administrator  
St. Peter’s Hospital

### Optimized protection for VMware infrastructure

Avamar deduplicates backup data globally, across physical and virtual servers. For virtualized environments, Avamar quickly and efficiently protects virtual machines (VMs) as if they were physical servers. Lightweight Avamar agents can quickly and efficiently deduplicate data at the VM Guest, Service Console, or VMware Consolidated Backup (VCB) proxy server. In all cases, only new, unique sub-file, variable-length data segments are transferred across the virtual/physical infrastructure for fast, daily full backups. Avamar can also deduplicate the data stored within virtual disks (\*.vmdk files), enabling efficient replication for disaster recovery. In addition, Avamar enables fast, single-step recovery of individual files or complete \*.vmdk files.

### Flexible solutions for remote offices

For smaller remote offices, lightweight, efficient Avamar software agents can be deployed on the systems to be protected with no additional remote hardware required. This allows data to be backed up directly over existing WAN connections to a central, cost-effective, and efficient Avamar server at the data center. Local tape backup and offsite tape shipments can be eliminated. To protect larger remote offices and provide faster recovery, data can be backed up to a local Avamar Data Store server or EMC-certified server and then replicated to another Avamar Data Store or certified Avamar server located at the disaster recovery site, hub site, or another data center if required.

**“The [EMC] Avamar remote office solution enabled us to reduce administrative support requirements by 80 percent, reduce backup windows by 90 percent, and recover lost files and servers in minutes rather than hours.”**

Mike DePhillip  
Backup Administrator  
Virginia DMV

### Relief for bandwidth-constrained data centers

Avamar also provides fast, daily, full backups for data center LAN and NAS servers. By deduplicating backup data at the source and globally, Avamar eliminates traditional backup bottlenecks caused by slow or congested networks, backup servers, and infrastructure. Avamar backup data can then be asynchronously replicated to a disaster recovery site for additional protection.

Avamar Feature	Avamar Benefit
Global, source deduplication	Reduces amount of data sent over networks and virtual infrastructure; reduces daily backup times by up to 10x, daily network impact by up to 500x, and cumulative backup disk storage by up to 50x
Wide area network (WAN) backups	Protects data with encrypted, electronic daily backups via existing networks; reduces or eliminates need for tape infrastructure
High availability and reliability	Fault tolerance across nodes with RAIN architecture for reliable data protection and access
Verified recoverability	Provides daily data recoverability and server integrity verification
Centralized management	Simplifies consistent, multi-site backup control operations from a single pane of glass and the automation of policy-based management
Fast, single-step recovery	Recovers data (whole backups, files, or directories) immediately; no need to restore last full and incremental backups
VMware infrastructure backups	Fast, efficient daily full backups via VM Guest, Service Console, or VMware Consolidated Backup (VCB) proxy server
Physical and virtual deployments	Optimizes solution for customer use environments

## Specifications

### Applications

DB2  
 IBM Lotus Domino  
 Microsoft Exchange  
 Microsoft SharePoint  
 Microsoft SQL Server  
 NDMP for NAS Filers (NetApp and EMC Celerra®)  
 Oracle  
 VMware

### Hardware

Prepackaged EMC Avamar Data Store configurations for simplified deployment; any certified Intel x86 server running Red Hat Enterprise Linux

### Operating systems

AIX  
 Free BSD  
 HP-UX  
 Linux  
 Mac OS X  
 Novell NetWare  
 Novell OES 2  
 SCO UNIX  
 Solaris  
 Windows

### Platforms supported

Please contact EMC for further details regarding certified hardware platforms.



EMC Corporation  
 Hopkinton  
 Massachusetts  
 01748-9103  
 1-508-435-1000  
 In North America 1-866-464-7381  
 www.EMC.com

### Take the next step

To learn how an EMC Avamar solution can make your backup and recovery processes more efficient and cost-effective, contact your local EMC sales representative or authorized value added reseller, call us at 1-866-464-7381, or visit our website at [www.EMC.com](http://www.EMC.com).