



EMC Backup and Recovery solutions for Virtual infrastructure

Aleksandar Antic
BURA – IP Technology Consultant
EMC Serbia

Content



- Challenges and solutions
- CDP
- VMware backup options
- Avamar
- Traditional backup solution
- Backup reporting

Virtualization – Data Protection challenges & solutions



- Challenges


- Resource utilization
- Disaster and operational recovery
- Duplicated data
- Management

- Solutions

- CDP
- DataDedupe
- Management tools

Old Paradigm

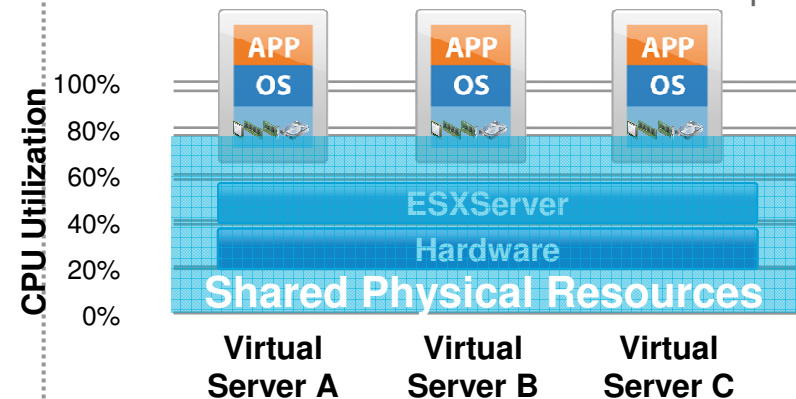
Physical Environment: Low overall server utilization and plenty of bandwidth for backup



20 percent resource utilization

New Paradigm

Virtual Environment: High overall server utilization and little bandwidth for backup



80 percent resource utilization

What is CDP?



Daily Backup: Recovery point every 24 hours



NetWorker / Avamar

Snapshot: Recovery point every n hours




SnapView / TF

Remote replication: Every I/O replicated, but susceptible to logical corruption



Mirrorview / SRDF

CDP: Application optimized recovery points

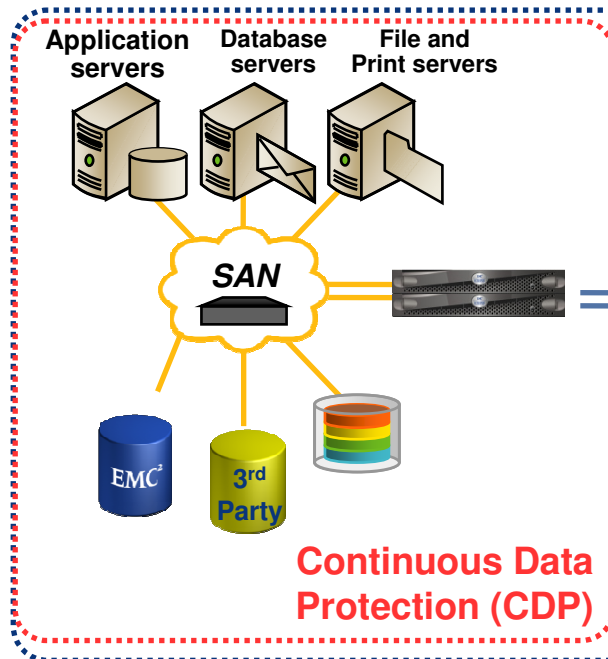


RecoverPoint

EMC RecoverPoint Overview

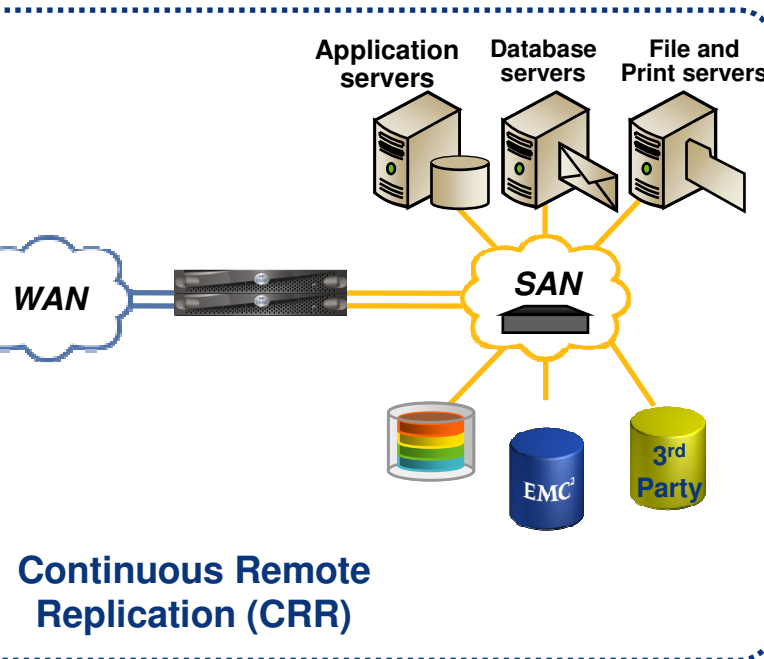


RecoverPoint CDP



- Utilizes host, fabric, or array write splitting
- Local replication with any point-in-time recovery
- Remote replication with point-in-time recovery
- Supports EMC and third-party storage
- Local and remote write journals enable roll-back to a consistent point-in-time image

RecoverPoint CRR and CLR

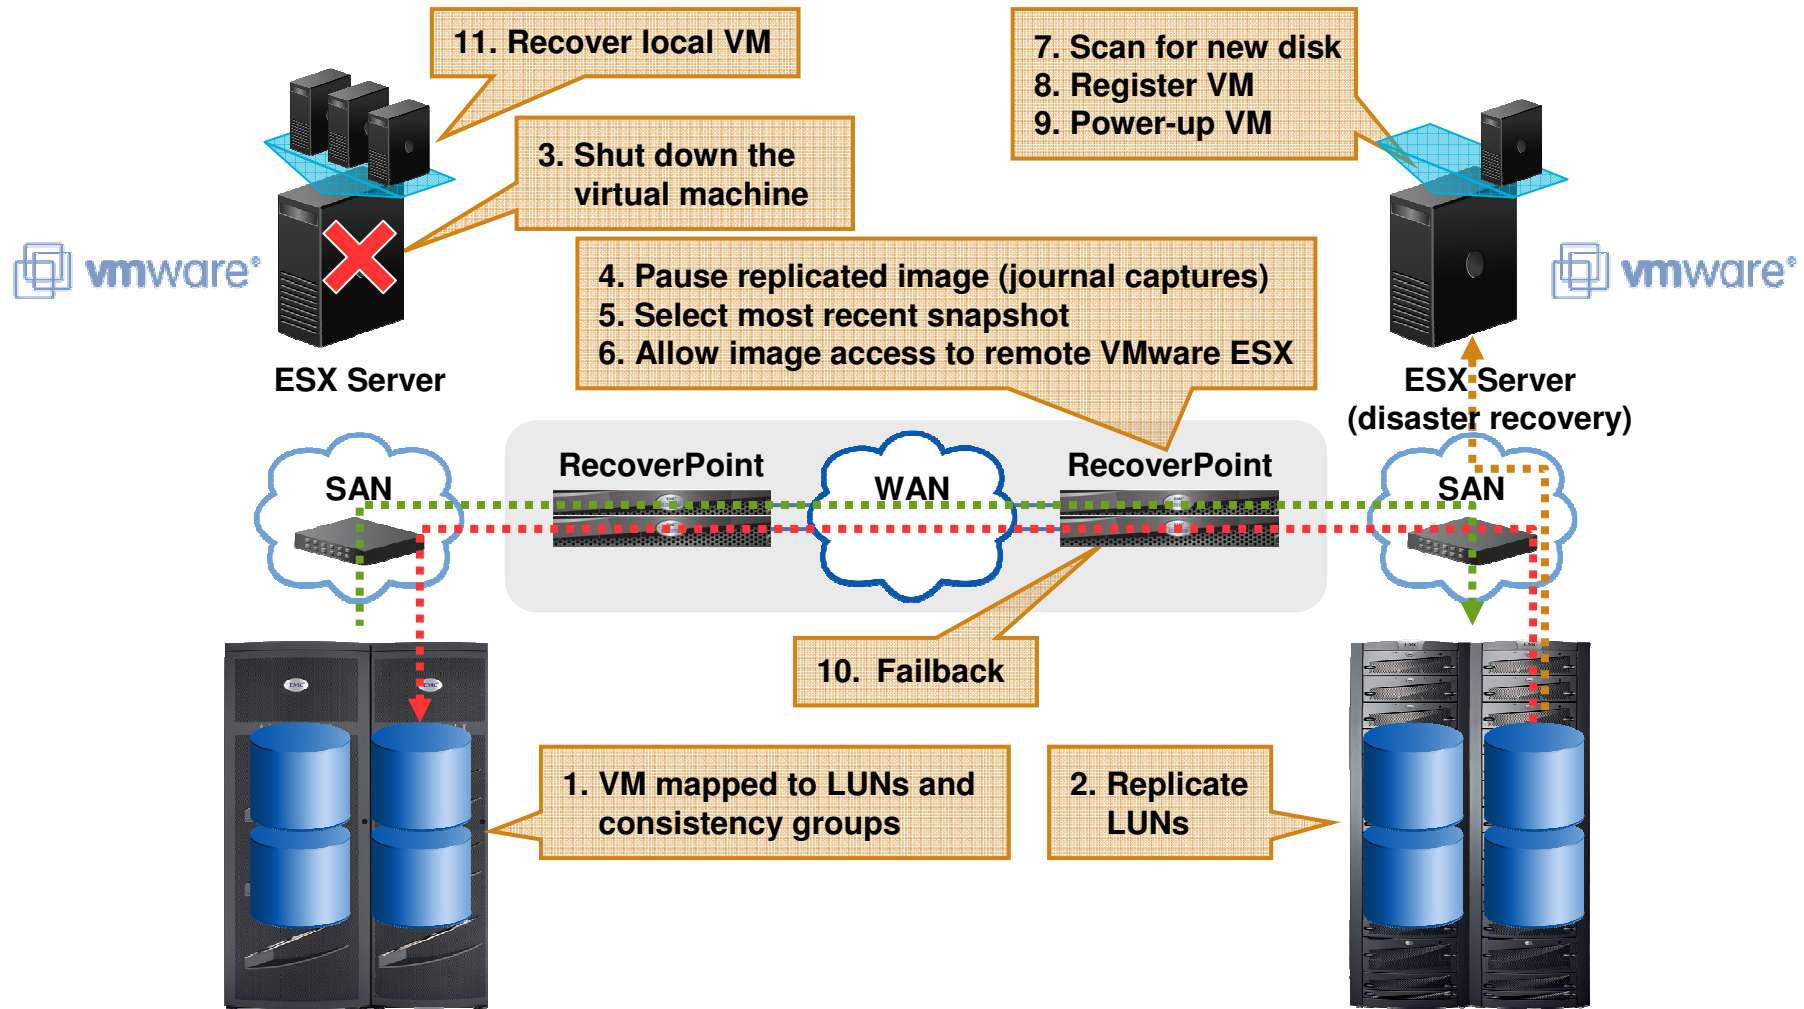


- Consistency groups ensure write-order fidelity within or across servers and virtual machines
- Map service level agreements using policy-based recovery point objective settings
- Block-level local protection for SAN, VMware Virtual Machine File System (VMFS), physical raw device mapping (RDM), SCSI/iSCSI (on CLARiiON CX4 or CX3 series volumes)
- Logical or physical access to replicated data


VM Failover **Before** VMware Site Recovery Manager



Manual Administrator Tasks




VM Failover with VMware Site Recovery Manager




ESX Server

Shut down the virtual machine
Pause replicated image (journal captures)
Select most recent snapshot
Allow image access to remote VMware
ESX Server
Scan for new disk
Register VM
Power-up VM
*** User initiates failback**
Reverse replication direction, sync data
Recover local VM



ESX Server
(disaster recovery)



SAN

RecoverPoint

WAN


RecoverPoint

SAN



1. VM mapped to LUNs and consistency groups

2. Replicate LUNs



Why RecoverPoint for VMware?



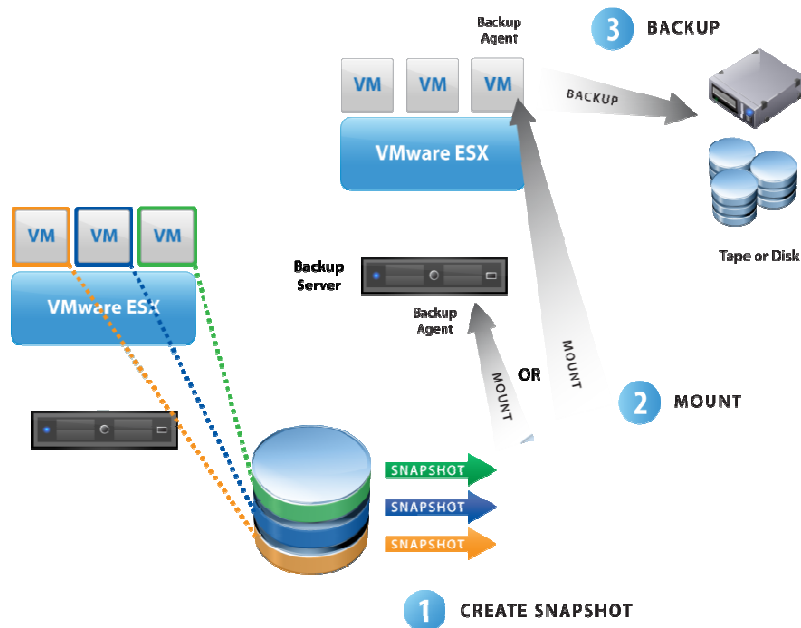
- Virtual machine-level continuous data protection against logical corruption
 - Roll individual virtual machines back in time
- Integration with VMware Site Recovery Manager
- Consistency groups for physical and virtual environments
- Local and remote replication in a single solution
 - Single point of integration with Site Recovery Manager
- Compression/data reduction minimizes communication costs
- Enables economical disaster recovery testing
- Application integration
 - Application consistent copies
 - VSS, VDI, Oracle
- Support for iSCSI and Fibre Channel-attached VMware ESX clusters

VMware backup options



- **Backup Agent in every VM** – Same as backup of physical machines (file-level backup)
- **vStorage APIs for Data Protection** (previously known as VMware Consolidated Backup -**VCB**) – vStorage API – image-level backup
- **Console backup** – not widely used because of limitations
- **VMware Data Recovery** – New tool integrated with vCenter Server to enable centralized and efficient management of backup of virtual machines.

vStorage APIs for Data Protection



Next Evolution of VCB

- Improved API enables native integration with partner backup application
- Deployable on Windows and Linux platforms
- Supports all storage architectures

Enhanced Functionality

- Supports incremental, differential and full VM image backup options
- Supports file level backup and restore
- Supports Windows and Linux guests

Customer Benefits

- Easy backup Integration with VMware vSphere
- Efficient backups
- Easy restore

EMC Solutions for VMware Backup



Avamar

- Complete backup, recovery, and dedupe solution
- Source/global deduplication at Guest, VCB, or Service Console
- Fast, daily full backups
- Single-step recovery
- Integrated high availability (RAIN)
- **Avamar Virtual Edition for VMware**



NetWorker

- Industry-leading backup and recovery software
- Integration with Avamar—and other EMC backup and data protection technologies



Data Protection Advisor

- Monitor data protection infrastructure
- Troubleshoot specific failures and issues
- Trend and analyze capacity, SLAs, more
- Plan by analyzing environment-wide usage
- Predefine, customizable reporting

EMC Avamar - How it Works 1/2




Global Source-based Data De-duplication

- ① Break data into atom
(sub-file, variable-length
segments of data)

- ② Send and store
each atom only
once

- ③ Avamar backup
repository



At the source—De-duplication before data is transported across the network

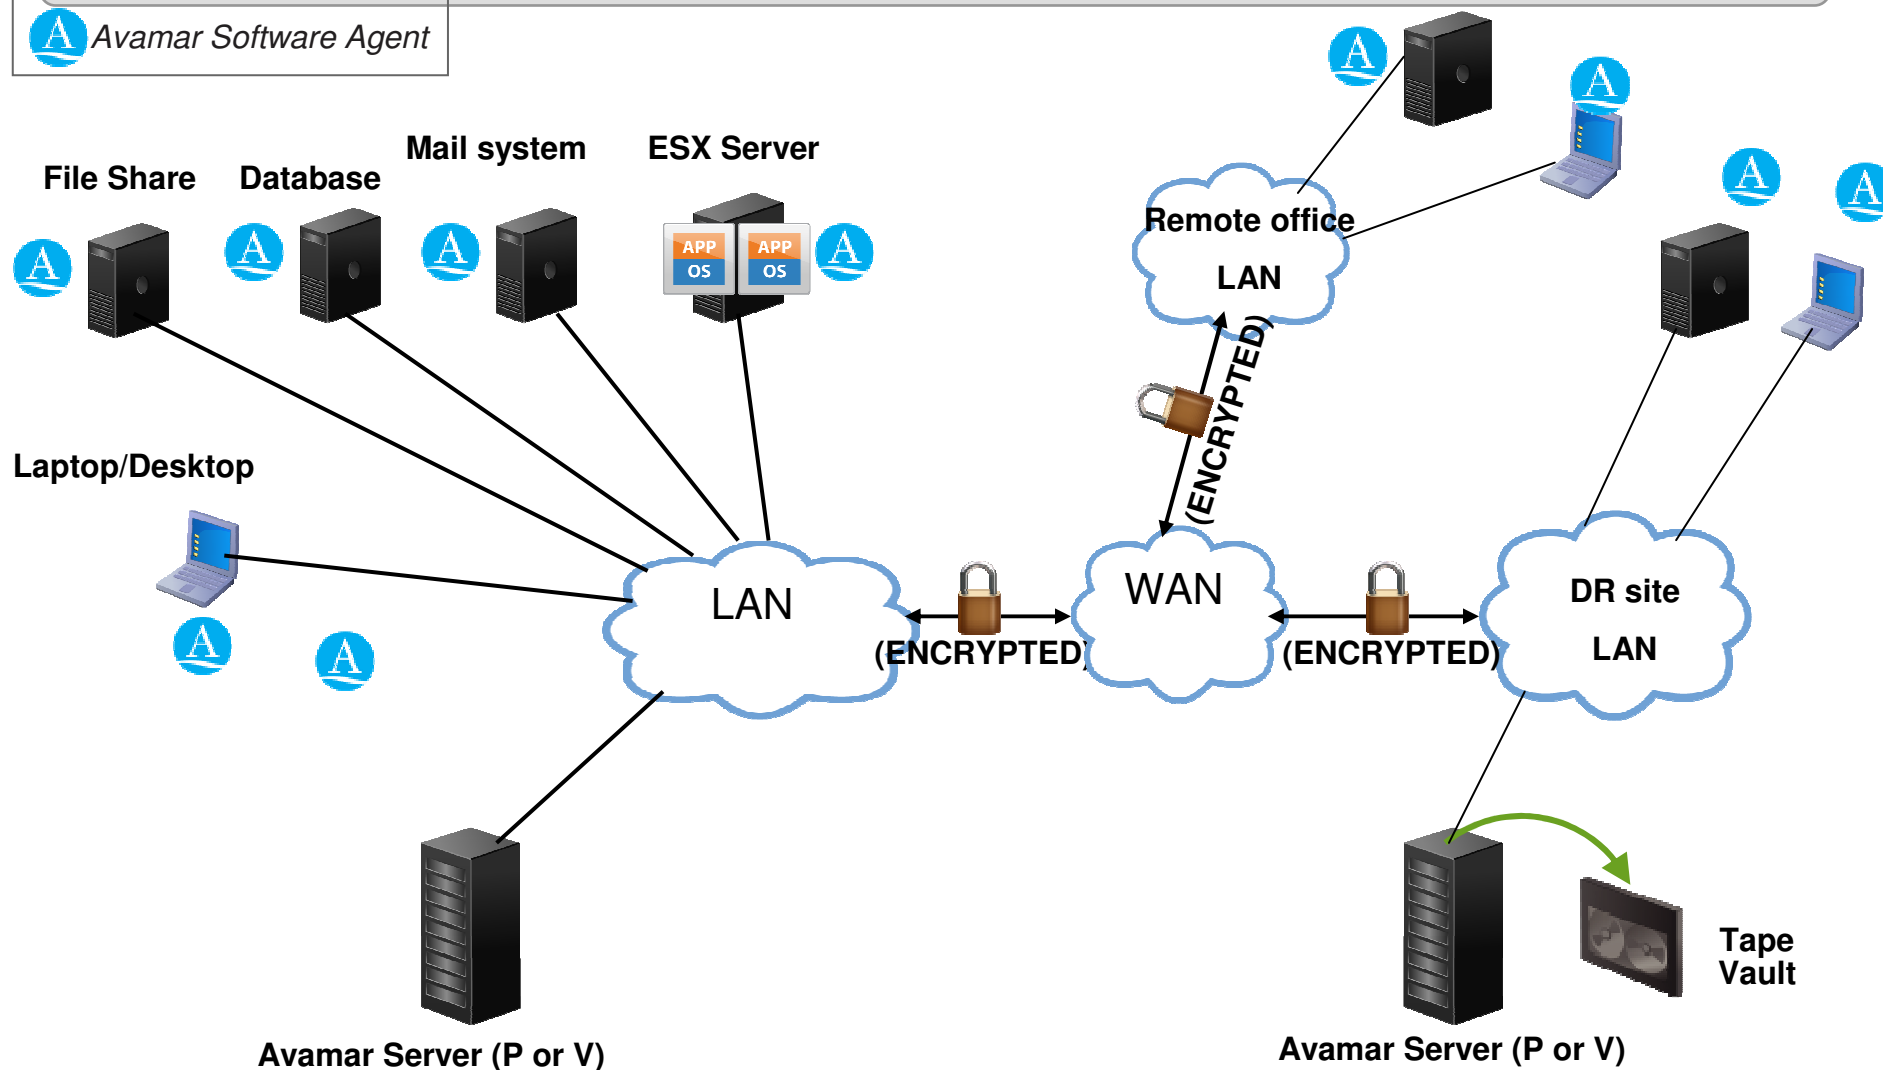
At the target—Assures coordinated de-duplication across sites, servers, and over time

Granular—Small, variable-length sub-file segments guarantee most effective de-duplication

Avamar – How it works 2/2




 Avamar Software Agent



Avamar Backup Solutions for VMware



VMware Guest OS Backup



VMware Virtualization Layer

x86 Architecture


= Avamar Software Agent

Avamar client software runs directly on each virtual machine

vStorage APIs for Data Protection

= Avamar Software Agent

Virtual Machines



Physical Server



SAN Storage

SNAPSHOT
SNAPSHOT
SNAPSHOT

Centralized Data Mover



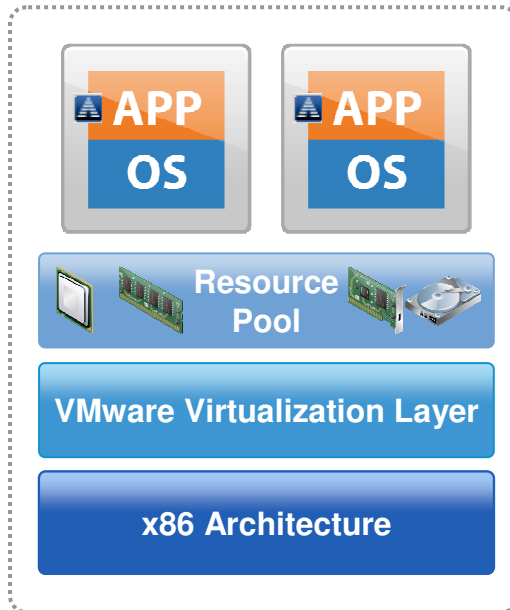
Avamar server



vStorage API proxy server with Avamar agent

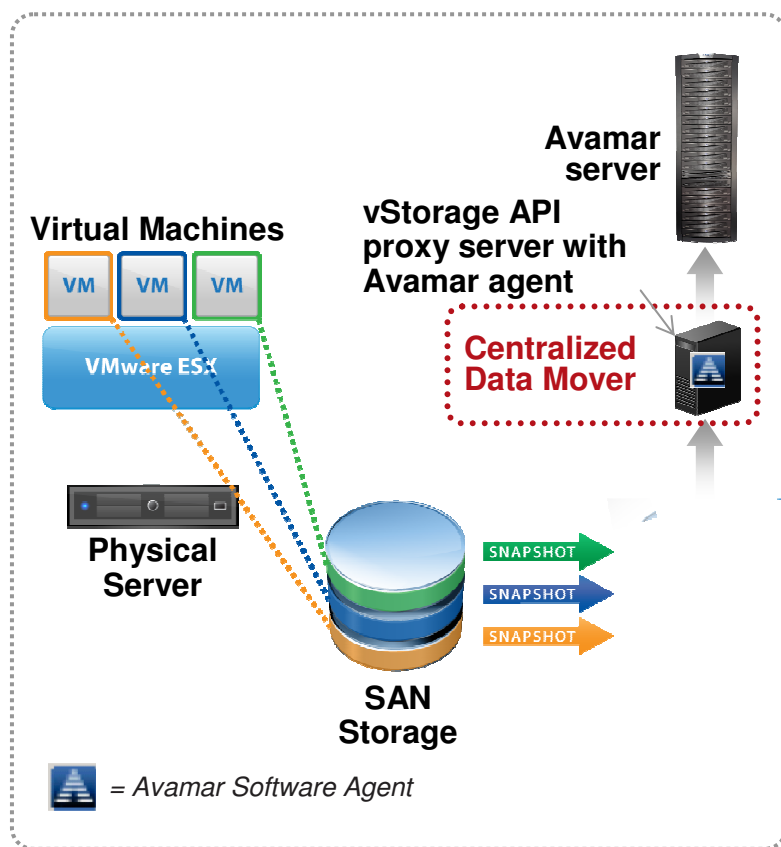
Avamar client software runs on the proxy server

Avamar Backup for VMware Guests



- Avamar agent resides inside each virtual machine
- Deduplicates data within the virtual machine, as if they were physical servers
- Moves minimal backup data
 - Reduces resource contention and accelerates backups
- Provides file-level restore for Windows, Linux, and Solaris

Avamar and vStorage APIs for Data Protection

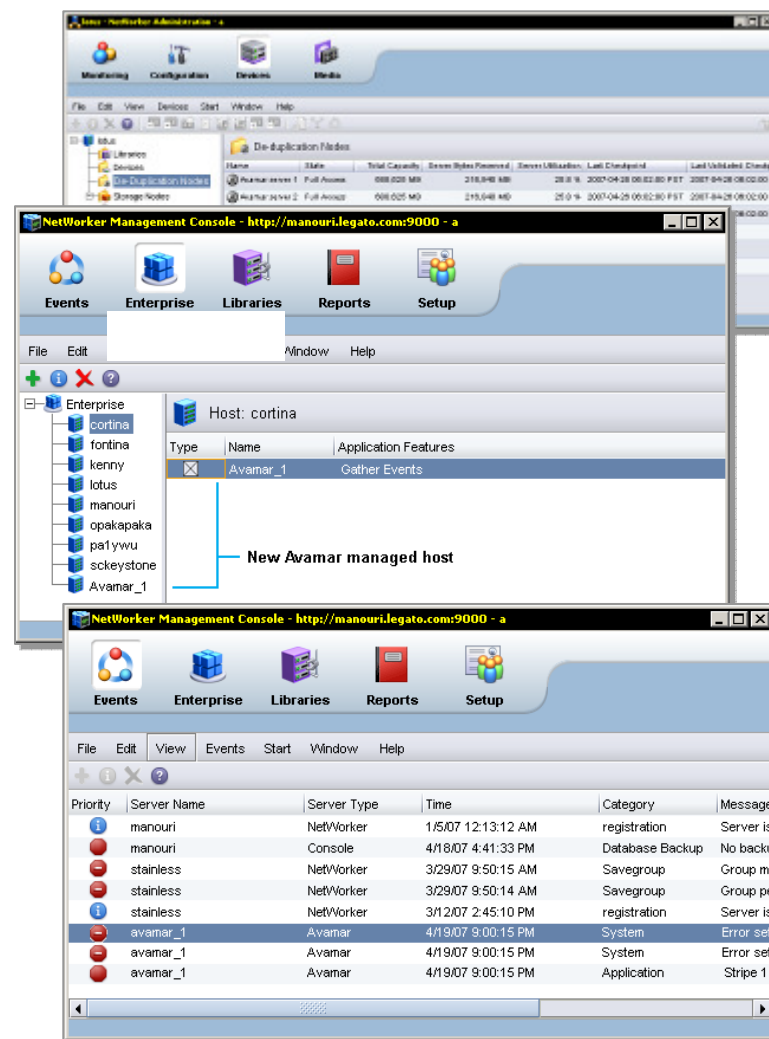


- Leverages new VMware vStorage API for data Protection
- Avamar agent hosted on the Proxy host
- Proxy Server is hosted on a VM
- No agents running on the target VMs
- By leveraging SCSI Hot Add, VMDK files are available to the Proxy
- Option to leverage change block feature (reduces file scan time)
- Can run either Windows or Linux Proxies
- Supports ESX 4.0
- 3 restore options:
 - Same VM
 - New VM
 - On top of existing VM

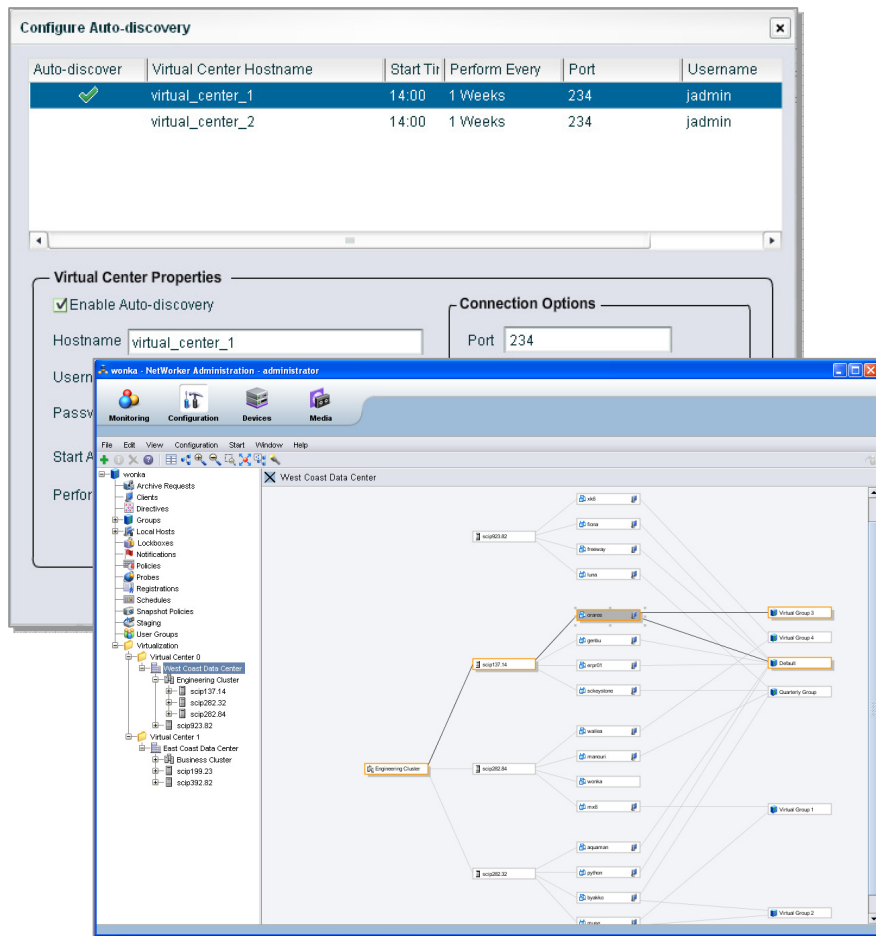
Traditional backup - EMC NetWorker



- Centralized control of traditional and next-generation backup
- Industry-leading source/global data deduplication
 - Reduces required total backup storage by up to 50x and data moved daily by up to 500x
- Broad backup to disk
 - Disk library integration, replication, snapshot management, and continuous data protection
 - NAS backup to disk
- Enterprise performance
 - Securely backups and reliable recoveries
- Better recoverability from tape backups
 - Future-proofed Open Tape Format with better recoverability from damaged tape media



VMware Auto-discovery and Visualization— Simplifying Management for VMware



- Auto-discovery

- Enables NetWorker to automatically discover new virtual machines added to your environment
- Establishes a connection between NetWorker Management Console and your existing VMware vCenter Server (formerly VirtualCenter) console
- Manages multiple virtual centers in one application
- Customize to run when you need it

- Visualization

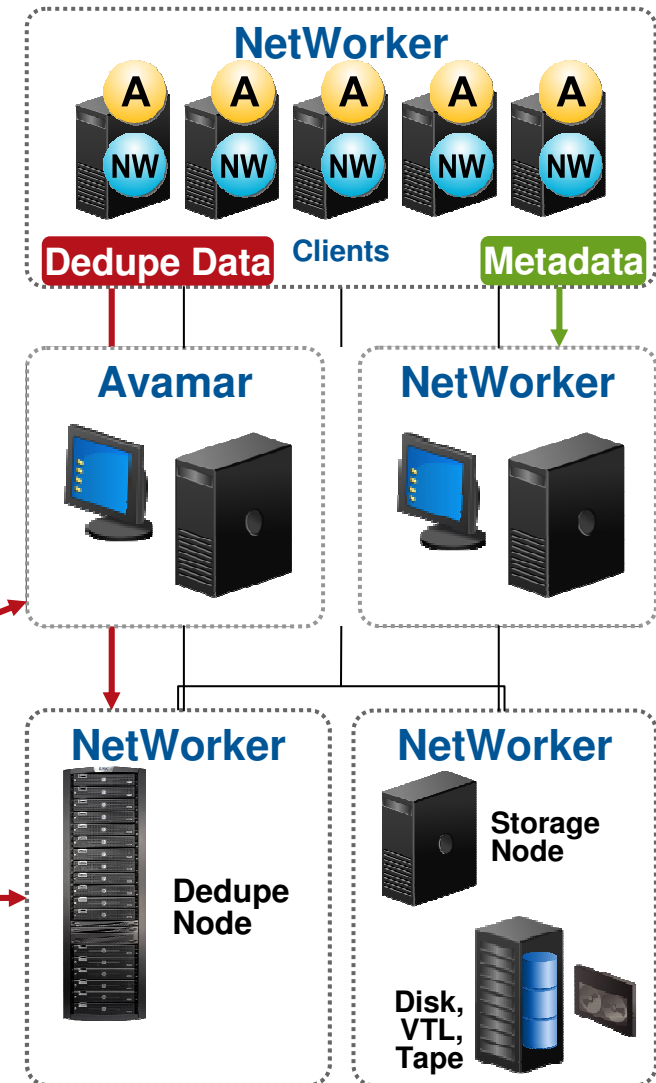
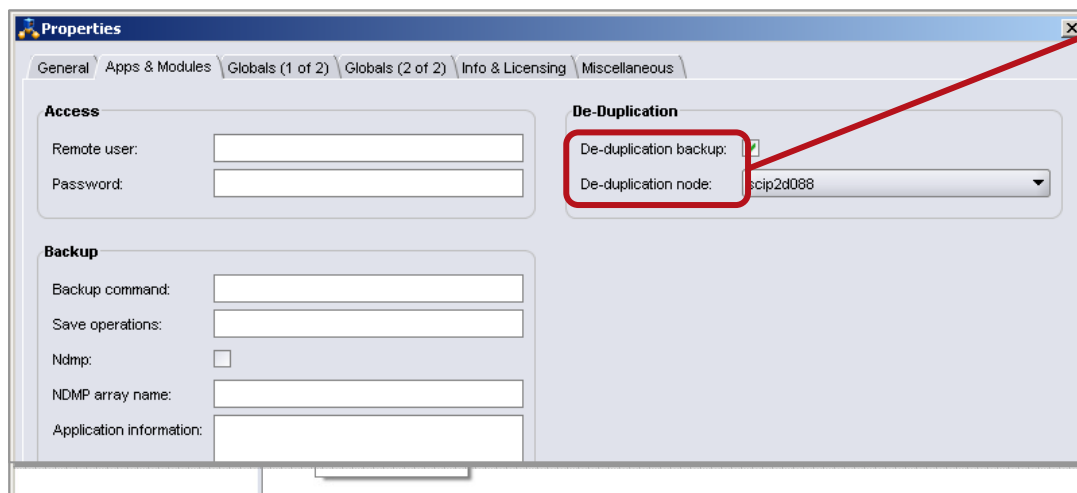
- Visual map of your virtual environment
- Displays relationships between elements in your virtual environment and NetWorker
- Interactive controls allow users to manipulate the map to filter contents
- Configure from map—right-click to launch wizard

EMC NetWorker and Deduplication

NetWorker and Avamar integration




- NetWorker client and Management Console communicate with Avamar
- Avamar appears as a NetWorker dedupe node enabled via client properties
- NetWorker manages work flow—metadata indexing and data sent to the dedupe node



Traditional backup



- Advantages:
 - traditional way
 - OS and applications coverage
- Disadvantages
 - traditional way
 - LAN utilization
 - Sometimes complex to configure and manage



Backup hardware



- Tape and tape libraries
 - Traditional destination for backups
 - Sequential access
 - No protection
 - No advanced features (like deduplication, replication from one to another library etc)
 - Sometimes difficult to configure, manage and troubleshoot
- Disk and disk libraries (Virtual Tape Libraries – VTLs)
 - Random access
 - Protected by the storage array (RAID, hot spare, etc)
 - possibility of having Data Deduplication (3D), replication and other advanced features
 - Easier to configure, manage and troubleshoot

Today, most of the customers is using combination of the disk and tape backup!

Inline Deduplication Systems



DD500-DD600 Series



DD690



DD880



DDX Array Series

Up to 16 Controllers

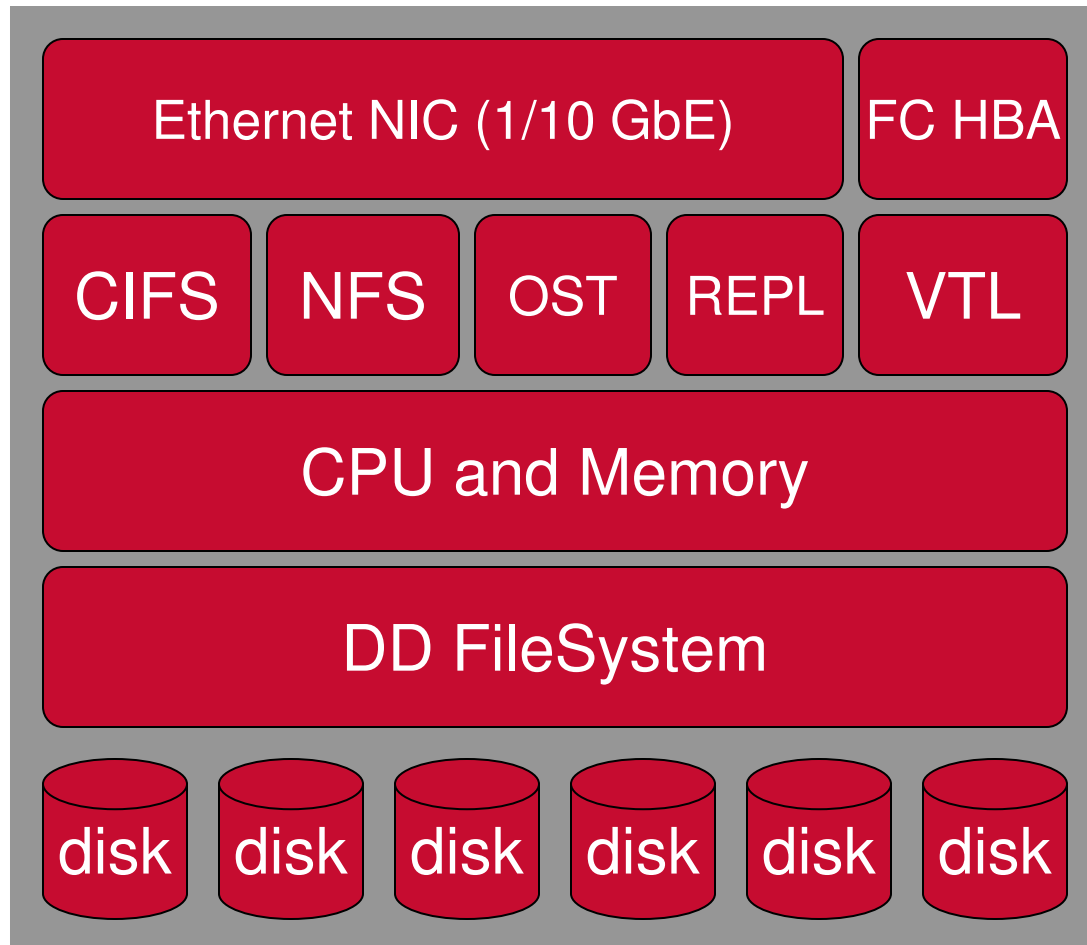


DD120

Remote Office System

	DD120	DD510	DD530	DD565	DD660	DD690	DD880
Speed (GB/hr)	300	435	540	1 TB/hr	2 TB/hr	2.7 TB/hr	5.4 TB/hr
Logical Cap. (TB)	7	55	110	320	520	710	1.420
Usable Cap. (TB)	.373	1.3 to 2.7	2.3 to 5.4	Up to 16.2	Up to 26	Up to 35.3	Up to 71

Data Domain architecture



Use which interface suits you best, or both


Use which protocol suits you best, or both

Inline checksum, comparison, compression and prefetching.

Data containers in optimized dedicated filesystem

Can be internal S-ATA or a LUN of an external SAN

Multi Purpose, multi protocol




EMC Data Protection Advisor




Industry-leading data protection management solution

- Collects data from across the infrastructure
- Provides a central console for monitoring, alerting, analysis, trending, and reporting
- Transforms disparate data into actionable business information
- Reporting that is easily customizable
- Lowers costs, reduces risk, and improving change management

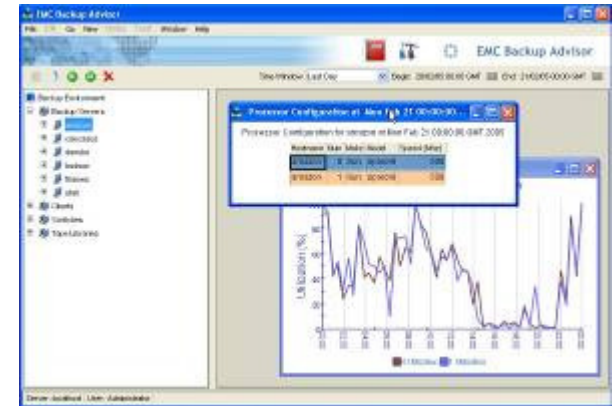
Be Proactive, Not Reactive



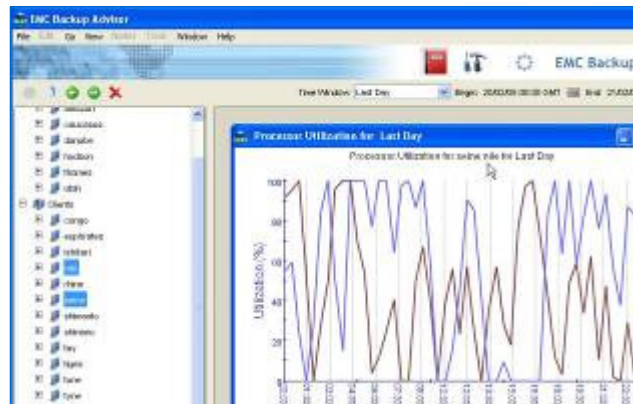
Reporting




Monitoring



Troubleshooting

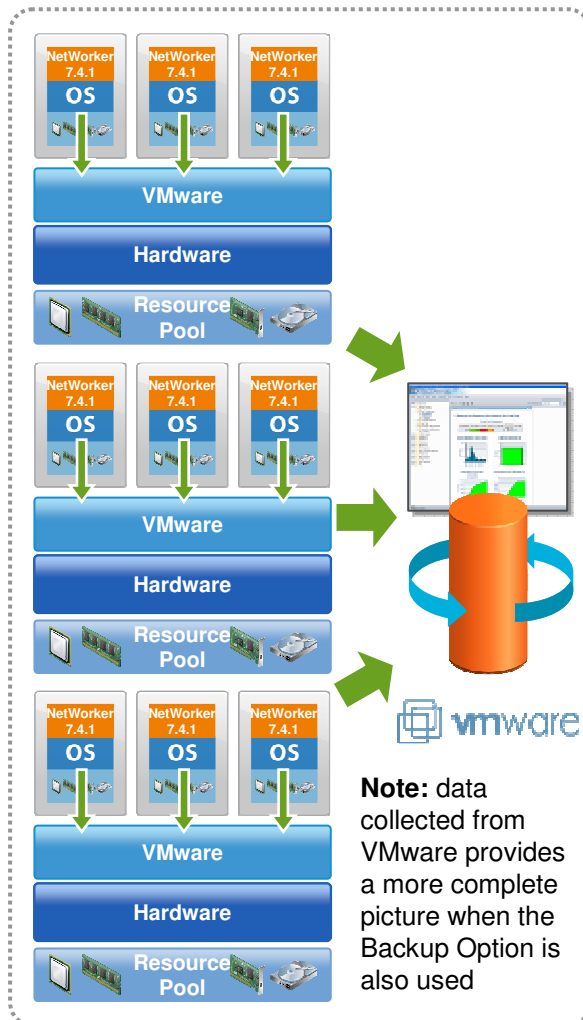


Performance Management



Capacity Planning

EMC Data Protection Advisor: Virtualization Option



Improve visibility and management for virtualized hosts

- Simplifies, centralizes management
 - Reduce the management complexity added by virtualization
 - Monitor, troubleshoot, optimize, and report across virtual infrastructure
- Collects data from vCenter or ESX servers

EMC²
where information lives[®]