PARAGON Software GmbH
Heinrich-von-Stephan-Str. 5c ● 79100 Freiburg, Germany
Tel. +49 (0) 761 59018201 ● Fax +49 (0) 761 59018130
Internet www.paragon-software.com ● Email sales@paragon-software.com

Paragon Protect & Restore (PPR)

User Guide

Contents

Introducing Protect & Restore

This chapter will help you get general information on a concept-new product from Paragon – Protect & Restore (PPR).

What is PPR?

PPR provides comprehensive agentless protection for virtual environments hosted by VMware vSphere or standalone ESX servers. It can also protect physical and virtual Windows systems remotely through on-site agents. Our solution supports a complete scenario for Microsoft Hyper-V that involves agent-based backup for the host and agentless for its guest machines. By operating at the application level PPR enables to create consistent backup images of MS Exchange databases without any impact on the production email server. As a backbone it uses Paragon's brand-new distributed architecture that allows efficient central management of hundreds of computers.

Being obviously a product for IT personnel, it's anyway really flexible in administration. For those, who are accustomed to automation scripts and command-line tools, there's a Windows PowerShell based console, while for those, who value comfort – a well thought-out GUI console.

Agentless Protection

More and more companies are adopting virtualization as a powerful tool for consolidating hardware and infrastructure. Reduced IT costs, increased availability of hardware and applications and improved overall service levels are only few benefits virtualization can offer. However, despite of its tremendous promise, it also leads to a significant increase in storage capacity due to virtual machine image files and associated metadata that virtual servers create and run. So **the number one fact**: Virtualized systems generate more data, thus they demand more resources for protection.

The world of virtual machines is arranged in such a way that you CAN treat each virtual environment running on a virtual host as pure physical, thus applying traditional backup and disaster recovery strategies for its protection. But it can only be considered as a forced temporary solution as you won't be able to get advantage from agentless backup and replication functions, snapshot and changed block tracking (CBT) mechanisms of your hypervisor, its fast transport system and many other useful features. So **the number two fact**: A backup tool designed specifically for virtual environments is a priori much more flexible and efficient.

For VMware guest machines

PPR offers fast and reliable agentless disk-imaging backup and replication of online and offline virtual machines resided on VMware vSphere or a standalone ESX server. The use of VMware CBT (Changed Block Tracking) and the patent-pending Paragon's ITE (Image Transfer Engine) ensure full and incremental backups or replicas are created with the minimum time and impact on ESX. Flexible retention policies, enhanced data processing methods, automatic exclusion of irrelevant data (page files, etc.), and the innovative pVHD format allow the optimal usage of backup storages. Employment of MS VSS (Volume Shadow Copy Service) when taking snapshots of Windows machines, guarantees data consistency. Introduction of Paragon's ProTran®, a unique data transport protocol and a two-tier backup storage infrastructure, open up further minimization of backup windows and network traffic for simultaneously made backups. Paragon's GoForSure technology enables to be pretty sure created policies will be completed with a success.

When time comes for disaster recovery, our product does it in compliance with the most aggressive RPOs (Recovery Point Objectives). You've got the option to restore a backup to the original or a new location according to a certain restore point. When restored to a new location the target virtual machine will be appropriately reconfigured during the process. For replicas the whole disaster recovery procedure comes just to its launch, which may take a couple of seconds, thus ensuring the maximum business continuity. To make sure replication goes on as expected you can test any replica machine.

For Hyper-V guest machines

Operating at the virtualization layer through MS VSS API, PPR can also agentlessly protect virtual machines hosted by Hyper-V (Linux, Windows, and other OSes supported by this hypervisor). The current version can only protect entire virtual machines (online and offline) registered on a local Hyper-V host, where our Backup Agent and Hyper-V Application plug-in are deployed. The Tray Application utility is used to manage backup and restore tasks, thus it should be installed on the host as well. Centralized management through PPR Console is not yet available.

A Hyper-V guest machine can be backed up to a local disk or network share as a pVHD, VHD, or VHDX virtual container. When attempting to create another image of the same machine, our product prompts to reduce backup storage footprint by taking advantage of the incremental imaging.

When time comes for disaster recovery, PPR enables to restore a previously backed up virtual machine to the original or a new location according to a certain time stamp. If using VHD/VHDX as target backup format, you can attach the image to an existing Hyper-V virtual machine and OS will be launched successfully.

Agent-based Protection

PPR offers agent-based backup of Windows physical machines. Any Windows OS computer on the network can be protected entirely, by separate volumes or drive letters. Almost all technologies used for the agentless protection of VMware ESX guests are also available for machines protected through on-site agents (MS VSS, Paragon's ITE, ProTran, GoForSure, pVHD).

By embedding a special plug-in, users can monitor backup activities on target machines through the system tray. Wake-on-LAN Assistant allows waking up remote target machines to do backup. There's no need to install it on all machines that require it – the administrator just picks one and it will automatically wake up all others that share the same subnet when needed.

The agent-based backup technique can also be applied for protecting guests hosted by non-commercial VMware ESX, where the VMware snapshot mechanism is unavailable, or for VMware fault-tolerant configurations that do not allow agentless protection. Moreover, it can help to protect Windows OS guests of any other hypervisor.

When time comes for disaster recovery, data volumes can be restored remotely, while system volumes or entire machines – on-site from a special WinPE media. The third generation of Paragon's Adaptive Restore technology guarantees successful bare metal recovery of Windows OS systems to dissimilar hardware by injecting required drivers and other actions crucial for this type of migration. Paragon's Recovery ID allows minimizing time and effort of restore – the administrator sets up a one shot restore policy in the console, assigning it a particular ID. The user only starts up the failed computer from the WinPE recovery media and enters the obtained ID, thus initiating the pre-configured restore operation.

MS Exchange Protection

PPR offers agent-based protection of Microsoft Exchange Server 2007/2010/2013 and its email databases. By operating at the application level through MS VSS (Volume Shadow Copy) API, administrators have no need to allocate time for backup windows, for PPR enables to create consistent database backups without any impact on the production email server.

As for MS Exchange, PPR supports all latest backup technologies, like incremental backup chains, data retention mechanisms, block-level data de-duplication, replica databases, etc. But its main advantage is in flexible restore – restore of all or certain databases to the original or new location, including restore to RSG/Recovery Database with the option to create a dialtone database to let users send or receive emails in the process. When restoring the latest backup in the chain, there's the option to replay transaction logs, thus achieving minimal possible data loss.

PPR also allows non-destructive restore of certain mailboxes. By default, their contents will be restored to the original location, provided none of the already existed email items are lost. If necessary, you can specify any mailbox and a folder where you'd like the restored data to be placed to.

If using PPR together with Paragon Granular Recovery (PEGR), you can easily connect a backup email database to MS Outlook to view and extract certain emails.

Key Benefits of PPR

Supported VMware Hypervisors

PPR supports VMware ESX/ESXi 4.xx, ESX/ESXi5.xx, vSphere 4.0 - 5.5 configurations and non-commercial VMware ESX.

Supported Hyper-V Hypervisors

PPR supports Microsoft Hyper-V of Windows 8.1, 2008 R1/R2, 2012 R1/R2.

Supported Virtual Machines

You can protect any Windows, Linux, or other OS guest supported by VMware or Hyper-V. Through the on-site backup agent you can also back up Windows virtual machines of any other hypervisor available on the market.

Supported Physical Machines

You can protect any Windows OS machine since Windows XP.

Supported Storages

You can configure and use local (on Backup Server), ESX (on ESX datastore), network (UNC), and FTP/SFTP storages. Besides you've got the option to set up fully independent disk pools that unite from one to an unlimited number of external disks. The disk pool concept of PPR implies that you can add to or remove from a pool any disk at any moment without the need of reconfiguring backup policies that use this pool as destination. Once the first disk has been filled with backup data, another disk is used and so on.

Distributed Architecture

PPR has a modular architecture that allows flexible scalability for environments of different size and configuration. A unified installation package includes all components of the product, which can be deployed all on one machine, or be spread over several machines according to particular needs. Some components require manual deployment on-site (Protect & Restore Server and Console), the others can either be installed centrally from Console (recommended) or manually on-site when having trouble with the remote mode.

In complex environments consisting of hundreds of virtualized servers and workstations, distribution of roles among several machines enables to considerably speed up backup, replication or disaster recovery operations. By setting up a two-tier backup storage infrastructure for instance, you take workload off the primary backup server. Besides you're allowed to have several management points to run tasks or monitor the operation progress from any machine on the net.

VM Backup

Unlike traditional backup tools designed to work with physical machines, PPR can operate at the virtualization layer and directly employ the VMware snapshot mechanism to do backups. It doesn't need an agent on a target virtual machine to create its point-in-time copy including its configuration, operating system, apps, etc. This approach significantly enhances the backup performance, while minimizing the load on target machines and the hypervisor during the process. Besides, there's no need to provide credentials for every guest to do backups.

One backup task can involve one or many virtual machines. By default, for every machine our product creates a full backup in a special proprietary format during the first run, then incremental updates according to a set timetable. It allows configuring general retention policies for backup storages or a particular policy for a certain backup task, specifying how long backups should be kept or the amount of space they can take. When time comes, all restore points beyond the set limit are merged with their full backup thus creating a new full backup. All backup images are being highly compressed during creation by using redundant data exclusion filters (OS page files, zero data blocks, etc.) and a pVHD backup format, which eases the backup storage requirements.

Operating at the virtualization layer through MS VSS API, PPR can also agentlessly protect virtual machines hosted by Hyper-V (Linux, Windows, and other OSes supported by this hypervisor). The current version can only protect entire virtual machines (online and offline) registered on a local Hyper-V host, where our Backup Agent and Hyper-V Application plug-in are deployed. The Tray Application utility is used to manage backup and restore tasks, thus it should be installed on the host as well. Centralized management through PPR Console is not yet available.

VM Replication

For high-availability virtual environments that run the first tier applications, PPR complements VM Backup with VM Replication. Replication provides the best RTO (Recovery Time Objective), for this technique implies creation of clones (replicas) of target machines on a certain ESX datastore and their registering on the host under different names. Replicas are stored uncompressed in their native format, thus they are ready-to-go at any moment. All changes since the initial full replica are written to VMware native snapshot files, acting as restore points, thus allowing the usage of the VMware revert-to-snapshot mechanism to further accelerate disaster recovery scenarios, providing for almost zero downtime operation. You can also define a retention policy for replicas, thus all snapshots that breach the set policy will be automatically collapsed.

Physical Backup

PPR enables to protect any physical Windows machine since Windows XP. Target machines should be added to the infrastructure directly from the console or on-site to embed a special agent that will interact with the infrastructure and accomplish backup tasks. It’s possible to create a special policy that will periodically check Active Directory OUs for new members to automatically add them to the infrastructure. When setting up a physical backup policy, you can specify as a backup object entire computers or separate volumes. By default, for every machine our product creates a full backup in a special proprietary format during the first run, then incremental updates according to a set timetable. It allows configuring general retention policies for backup storages or a particular policy for a certain backup task, specifying how long backups should be kept or the amount of space they can take. When time comes, all restore points beyond the set limit are merged with their full backup thus creating a new full backup. All backup images are being highly compressed during creation by using redundant data exclusion filters (OS page files, zero data blocks, etc.) and a pVHD backup format, which eases the backup storage requirements.

MS Exchange Backup

PPR can treat MS Exchange 2007/2010/2013 at the application level, opening up the option to create consistent database backups without any impact on the production email server.

Incremental Imaging

When scheduling a backup or replica task, PPR employs a similar approach: For the first run it creates a full backup or a full replica, and then only saves changes since the last performed operation, increments for backup, and snapshots for replica. The delta to write is parsed through VMware CBT or Paragon’s ITE or both. This enables to maintain many restore points with the minimal backup storage requirements. Besides incremental updates require much less time to create, which takes the load off the whole infrastructure.

Enhanced Backup Format

Paragon introduces a pVHD (Paragon Virtual Hard Drive) format – a special VHD, optimized for storing backups of virtual and physical machines. It's very efficient in handling incremental chains, data de-duplication and synchronization. pVHD does away with all limitations of a standard virtual drive format, such as a poor compression ratio, integrity control, and encryption capabilities. As a result it allows obtaining backups that are up to four times smaller than original backup objects. If necessary, pVHD can be easily converted back to VHD.

Pre- and Post-scripts

To get consistent snapshots of virtual machines hosting applications that do not support Microsoft VSS, PPR enables to set for execution custom scripts before creating a snapshot to properly freeze these applications and after it to bring them back to normal work. PPR supports all popular script formats of Windows and Linux environments (.cmd, .bash, .sh, .tcsh, .bat, .php, .js, .vbs).

Backup Data Deduplication

Paragon's deduplication mechanism significantly reduces backup storage footprint. By linking existing local and/or network backup storages to Deduplication Server, you can make sure backup storages do not contain data duplicates, thus significantly cutting on storage requirements. Besides, it enables to cut even more on network traffic, as again only unique blocks of data are transferred to storages, thus having a positive effect on backup timings, RTOs and RPOs of company's IT infrastructure.

All backup data generated by various protection policies and sent to backup storages that are linked with one Deduplication Server will be deduplicated. Paragon's deduplication works the same way for both, virtual and physical backup images, thus enabling to achieve better deduplication efficiency. In combination with Paragon's dual backup mechanism, backup data deduplication becomes efficient for near-CDP scenarios (Post-processing deduplication) when the inline deduplication is either impossible or considerably slows down performance of incremental imaging.

PPR offers an intelligent continuous validation technique to guarantee all unique data blocks used by multiple restore points and backup storages are consistent. An additional level of protection can be achieved by using multiple deduplication servers in a mirrored configuration.

Two-tier Backup Storage Infrastructure

PPR supports a two-tier backup storage infrastructure that provides for further minimization of backup windows and network traffic for simultaneously made backups. In this type of infrastructure, you can allow the first-tier (primary) storage to reside as closer to target machines as possible, thus ensuring the highest backup or replication performance, while the second-tier (secondary) storage to be offsite (FTP, tape repositories, etc.), even on another continent, but huge and reliable. So during the dual protection process, first all target machines are quickly backed up or replicated to the primary storage, thus minimizing the impact on the production environment, and then these objects are being copied (archived) to the secondary storage during the night or a weekend as scheduled. Besides before copying all replicas are additionally converted to backups and vice versa to minimize the backup storage and network bandwidth requirements.

Backup Storage Browser

For easier administration of backups and replicas, all backup storages are open for browsing. Having a list of all created backups and replicas at hand, you can easily find and initiate an integrity checkup for those you consider critical, or delete those you don't need any more. When deleting an increment, you don't corrupt the whole chain, for our product automatically reorganize it for you.

Backup/Replica Validation

PPR enables to create rather complicated backup and replication policies that include a lot of parameters. To make sure a task will be successfully completed, there's a special mode (Save & Validate), when the corresponding backup or replication policy doesn't run till the end, but a particular step, which is defined by a level of validation (fast, medium, thorough). The "Fast level" includes checkup of all policy rules and their parameters, availability of the storage and ESX connection parameters. The "Medium level" includes connection to the specified ESX host to scan for target virtual machines as well as connection to the storage to retrieve metadata from it. The "Thorough level" includes creation/deletion of snapshots of target virtual machines, creation of an uncompleted backup/replication session and data items in the storage without opening data streams and data copying. The Save & Validate mode is always active for backup and replication tasks.

VM Recovery

With PPR you can recover a virtual machine to any good-to-know point in time and place it to the original or a new location. When restored to the original location, the original machine will be deleted (it should be offline). When restored to a new location you will be prompted to provide a new name for the machine, a host and datastore to reside it, virtual disk type, and network properties. Our product will change the VM configuration file and store the target machine according to the defined location. For replicas the whole disaster recovery procedure comes just to its launch (replica failover), which may take only a couple of seconds.

Replica Test Failover

PPR helps to test the sanity of any time stamp of an existing replica machine by non-disruptively simulating recovery procedure in an isolated network environment to make sure a certain replication policy produces valid replica machines, do field test for an existing recovery plan to rely on it in case of disaster or train IT personnel on what is to be done in case of emergency.

In PPR, backup techniques for physical and virtual machines are closely integrated, which opens a way to easily convert backup images of physical machines to VM replicas as part of a dual backup process that involves archiving to ESX-based secondary storage. Beside other advantages, obtaining ESX replica machines out of physical backup images allows Administrator to benefit from the replica failover / replica test failover functions.

Physical Recovery

With PPR you can recover a physical machine to any good-to-know point in time. Data volumes (non-system volumes) can be restored remotely, while system volumes and entire machines – on-site with a special WinPE media prepared beforehand through Recovery Media Builder. To do a restore operation the target computer should have a network connection to Administration Server or one of Backup Servers. Thanks to the third generation of Paragon's Adaptive Restore technology, the same hardware components are no more a demand – you can restore a Windows based system (any since Windows XP) to a completely different hardware by injecting required drivers and other actions crucial for this type of migration. Paragon's Recovery ID allows minimizing time and effort of restore operations – the administrator sets up a one shot restore policy in the console, assigning it a particular ID. The user only starts up the failed computer from the WinPE recovery media and enters the obtained ID, thus initiating the pre-configured restore operation.

Launch Backup

The launch backup aka instant restore is another feature that helps you minimize downtime of a failed production system. It enables to immediately run a physical or virtual machine directly from one of available restore points in VMware ESX environment. Thus users may continue their activities, while you've got enough time to pinpoint and fix the failed system.

File-Level Recovery

PPR allows browsing contents of virtual or physical backup images as well as VM replicas to do granular recovery of separate files and/or folders. Required data can be restored either locally (on a machine where Protect & Restore Console is installed) or on a network share, provided the original directory structure is kept intact if necessary.

MS Exchange Recovery

PPR enables to recover Microsoft Exchange Server 2007/2010/2013 at the database or mailbox level. Besides it can be used together with Paragon Granular Recovery (PEGR) to recover certain email data directly in MS Outlook.

Separate data stores or storage groups or all groups at once can either be restored to the original or some alternative location with the option to create a dialtone database if necessary. It's also available the so-called non-destructive restore of certain mailboxes. By default, their contents will be restored to the original location, provided none of the already existed email items are lost. If necessary, you can specify any mailbox and a folder to place the restored data to.

Role-based PPR Security

Only authorized users are able to access the PPR infrastructure. It is realized through a special authentication handshake. The entire authorization process is administered by Authorization Server, which includes the "token service" and a plug-in to interact with external authorization facilities (MS Active Directory or Windows Workgroup). All traffic after a successful handshake is encrypted, eliminating any chance of confidential data leak.

Thanks to a role-based security model, the PPR Administrator can effectively manage user privileges. There are several predefined security groups that are automatically created during the product installation (locally for a workgroup environment and in Active Directory for a domain). Users of each group have a certain set of permissions (operations they are allowed to accomplish in PPR). In general, PPR fully complies with the Bronze Certification requirements (at least two security groups, traffic encryption, support of user account policies, like password complexity, expire date, etc.).

Local and network backup storages can also be encrypted against unauthorized access through an industry-standard 256-bit AES encryption algorithm.

Realtime Statistics, Notifications, and Reporting

PPR includes smart tools of monitoring. You will be notified through popup windows on start of run-now policies. In Console you can get real-time statistics on any activity executed at the moment, filter error, warning or information infrastructure events, set up notifications and reports by a particular type of events or several types (errors, warnings, information, etc.) for an infrastructure component or activity you're interested in (backup or replica policies, etc.).

Our company is constantly releasing new versions and updates to its software, that's why images shown in this manual may be different from what you see on your screen.

How PPR Works

Terms

Term	Definition
Administration Server	The main service of the infrastructure. It manages all members of the infrastructure and controls usage of shared resources. It also includes Installation Server at the moment
Installation Server	The main service of remote installation and update of infrastructure components
Backup Server	A service that takes care of storage and maintenance of backups and replicas
Deduplication Server	A service that controls deduplication of backup data
ESX Agent	A service that interacts with the VMware infrastructure and does backup or replication of ESX guests
Agent	A client service that conducts maintenance of a single machine
Backup Agent	A client service that conducts backup of a single machine
Plug-in	A module that extends functionality of agents
Console	Provides user interface for management and control
Role	Purpose of a particular functionally independent group of modules
Policy	A synonym for a task or job involving one or many participants
Rule	Any parameter of a policy is a rule, except for its type
Task	Every launch of a policy specific operation on a particular machine
Backup Item	A particular backup object, e.g. database, volume, file
Backup Storage	Resource for storing backup images and VM replicas
Backup Catalog	Management structure on Backup Storage that keeps information on backup images
Backup Session	Any task that has to do with a backup image, either created or stored on Backup Server
Installation Client	An agent of remote installation
Database Replication	The process of synchronizing local databases with the main database
Repository	Resource for storing installation packages (a local folder on Installation Server)
Events	Any event that takes place in the infrastructure (collected from all servers and agents)
Event Viewer	A utility to view infrastructure events

Architecture

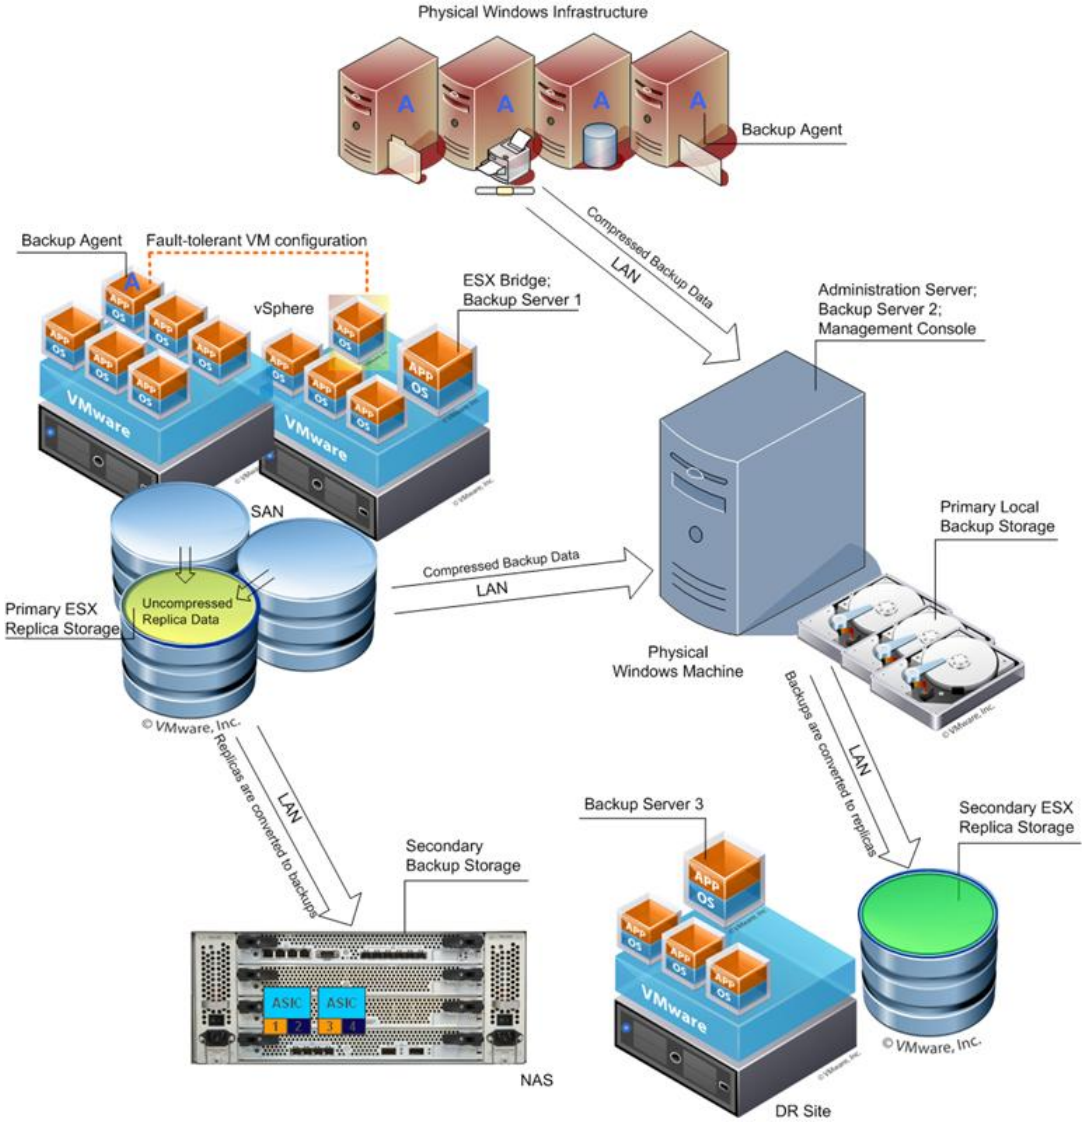

Policies

In PPR most actions are done by submitting corresponding policies. There are service and user policies. Service policies are created automatically by an event, for instance when adding new target machines to the infrastructure. The administrator can only create and manage user policies.

A user policy determines:

- Operations to execute
- Their parameters
- Executers
- Schedule.

Rules

In PPR each policy is presented by several policy objects. The main properties of a policy are called rules, which in their turn are presented by rule objects. For instance, an operation schedule is a rule, the same goes for a backup item (object), and so on.

Roles

In the PPR infrastructure each machine acquires one or several roles. Roles describe functions of components installed on a given machine. Components are grouped by functional identifiers. Each group forms a particular functionally independent service, e.g. Installation Server, Backup Agent, Backup Server, etc. Thus these groups of components are called roles. Components of each role are always installed and set up together, in order to work as a single fully functional service.

There can be the following roles:

- **Administration Server** (the main service of the infrastructure that keeps the central database, distributes resources, solves conflicts between the others, etc.)
- **Installation Server** (conducts deployment or update of components)
- **Backup Server** (responsible for storing and managing backups and replicas on local, ESX, UNC (network), and FTP storages)
- **Deduplication Server** (conducts deduplication of backup data)
- **ESX Agent** (interacts with the VMware infrastructure and does backup or replication of ESX guests)
- **Agent** (executes operations on a single machine)
- **Backup Agent** (conducts backup of a single machine)
- **Wake-on-LAN assistant** (allows waking up remote target machines to do backup).

It's up to the administrator to choose a role for this or that computer. One computer can have several roles at a time, but it's not allowed to have several Administration Servers in the PPR infrastructure. Roles can be changed from the Console.

Administration and Installation Servers are always installed on one computer having a general name of Protect & Restore Server in the Installer.

Administration Server

It's the main service of the infrastructure that has the following functions:

- Maintenance of the central database
- Synchronization of the central database with local databases
- Commitment of policies
- General operation planning
- Distribution of resources
- Primary operation analysis and conflicts solving.

Administration and Installation Servers are always installed on one computer having a general name of Protect & Restore Server in the Installer.

Installation Server

It is a service of remote installation and update of PPR components. Its main functions are:

- Installation, removal, update of components on all machines that join the infrastructure
- Management of the repository that contains all installation packages

All actions of Installation Server are presented by different installation policies. The administrator cannot directly create or modify this type of policies, they are formed automatically by an event, for instance, when changing roles assigned to a particular machine.

Administration and Installation Servers are always installed on one computer having a general name of Protect & Restore Server in the Installer.

Backup Server

It's a service that is responsible for storing backups and replicas on storages registered on this server. PPR supports the use of several backup servers, besides on each server you can register several primary and secondary storages.

Deduplication Server

It's a service that is responsible for deduplicating backup data on primary and secondary local/network backup storages. PPR supports the use of several deduplication servers.

ESX Agent

It's a service that interacts with and employs facilities of the VMware infrastructure to do backup or replication of ESX guest machines, thus providing for much higher performance and safety of backup or replication operations. For the maximum performance, ESX Agent should be installed on a Windows OS guest of an ESX host that has access to datastores with virtual drives you're planning to protect. Its installation can be initiated either from the Console (default way) or manually on-site.

Agent

It's a client service that executes service operations on a single machine that joins the infrastructure. By creating a policy in the Console, the administrator initiates one or several agents. Its installation can be initiated from the Console.

Backup Agent

It's a client service that protects data of a single machine at various levels (sector, file, or business apps). Its main functions are:

- Analyzing possible backup items of the given machine
- Execution of operations required to correctly back up protected backup items
- Splitting backup streams by backup items
- Collecting information on the protected machine, configuration of business applications
- Plug-in system to extend functionality
- Offline operation after receiving tasks.

Wake-on-LAN Assistant

It a service that allows waking up remote target machines to do backup. There's no need to install this plug-in on all machines that require wake-up-on-LAN. Just select one that is always online and it will automatically wake up all others that share the same subnet when needed.

Console

PPR actually includes two consoles, one GUI and one Windows PowerShell based. Both serve for administering the infrastructure. They don't acquire any role in the infrastructure and can be installed on several computers on the net, providing for the so called dynamic site management.

Both consoles are always installed together. In this guide only the GUI console will be considered for administering the infrastructure.

Getting Started with PPR

In this chapter you will find all the information necessary to get the product ready to use.

System Requirements

PPR Core Components: Administration/Installation Server, Backup Server, Management Consoles	
Operating System	- Microsoft® Windows® XP SP3 - Windows Vista - Windows 7 - Windows 8 - Windows 8.1 - Windows Server® 2003 SP2 - Windows Server 2008 R1/R2 - Windows Server 2012 R1/R2 **Note!** VMware Advanced Transport is only supported under Windows Server 2003 SP2/2008/2012
Architecture	32/64-bit
Memory	2GB or higher
Processor	x86/x64 processor (minimum 2 cores / 4 cores for Backup Server). The use of fast multi-core processors improves the data processing performance, and allows for more concurrent jobs
Storage	- Administration/Installation Server: 400 MB to install - Backup Server: 200 MB to install - Console: 50 MB to install **Note!** If you'd like to use disk pool (rotating media) as storage tie, please make sure Backup Server is installed on a physical machine
Environment	Active Directory domain environment or Workgroup environment
Credentials	- Domain administrator credentials - Domain administrator must have permissions of the local domain administrator on all machines that join PPR infrastructure
Services	- Running WMI service (Windows Management Instrumentation) - File and Printer Sharing is allowed in Firewall, which enables default system shares on local volumes (ADMIN$, C$, D$, etc.)

PPR ESX Agent	
Operating System	- Windows Server® 2003 SP2

	- Windows Server 2008 R1/R2 - Windows Server 2012 R1/R2 - Windows Hyper-V Server 2012
Architecture	64-bit
Memory	2GB or higher
Processor	x64 processor (minimum 2 cores). The use of fast multi-core processors improves the data processing performance, and allows for more concurrent jobs
Hypervisors	- VMware® vSphere® 5.5, 5.1, 5.0, 4.1, and 4.0 (any guest OS supported by VMware) - VMware ESX® and VMware ESXi™ 4.x and 5.x (any guest OS supported by VMware) - Non-commercial VMware ESX (Windows OS guests only through on-site backup agents) - VMware fault-tolerant configurations (Windows OS guests only through on-site backup agents) - Other hypervisors (Windows OS guests only through on-site backup agents)

PPR Agents	
Operating System	- Microsoft® Windows® XP SP3 - Windows Vista - Windows 7 - Windows 8 - Windows 8.1 - Windows Server® 2003 SP2 - Windows Server 2008 R1/R2 - Windows Server 2012 R1/R2 - Windows Hyper-V Server 2012
Architecture	32/64-bit
Memory	2GB or higher
Exchange Servers	- Exchange Server 2007 - Exchange Server 2010 - Exchange Server 2013 - Single Copy Clusters (SCC) support - Cluster Continuous Replication (CCR) support

If the 'Windows Firewall/Internet Connection Sharing (ICS)' service had been stopped during the installation, and then started by the administrator, Protect & Restore Console would not be able to connect to Administration Server. To tackle this problem, please

consult the PPR and Windows Firewall chapter.

Infrastructure Deployment

Deployment Phases

Mandatory actions on-site

To start working with PPR it's enough to install two components, namely **PPR Server** and **PPR Console**. You can do it through the PPR Installer. Once done, you can launch PPR Console to remotely install and manage all other infrastructure components.

Getting the infrastructure ready to work

Depending on machines and applications you're planning to protect and preferred backup storage tiers, the following additional actions might be required:

- **Installation of Protect & Restore ESX Agent** to protect guests of one or several ESX hosts or vSphere.
- **Registering of a standalone ESX host or vCenter** to access its guests for protection and/or management.
- **Installation of Protect & Restore Backup Server** to configure backup and/or replica storages.
- **Setup of backup and/or replica storages**.
- **Installation of the Backup Virtualizer plug-in** to accomplish launch backup (instant restore) tasks.
- **Installation of Protect & Restore Deduplication Server** to deduplicate backup data.
- **Setup of deduplication storage**.
- **Adding Windows physical machines to the infrastructure** to do agent based protection.
- **Installation of the Exchange Server 2010/2013 application plug-in** that extends functionality of the general agent allowing it to back up and restore databases of MS Exchange Server 2010/2013.
- **Installation of the Exchange Server 2007 application plug-in** that extends functionality of the general agent allowing it to back up and restore databases of MS Exchange Server 2007.
- **Installation of the Hyper-V Application plug-in** that extends functionality of the general agent allowing it to back up and restore guest machines resided on the local Hyper-V host. Please note that the current version of the product supports the mentioned scenarios through the Tray Application utility only.

All the mentioned above components can be installed centrally through PPR Console (recommended) or manually on-site through a custom mode of the PPR Installer.

Additional actions

- **Creation of WinPE recovery media** to do bare metal recovery of Windows physical machines. It can be done with Boot Media Builder.
- **Installation of Protect & Restore Tray Application** on a Windows physical or a virtual machine protected by Backup Agent. The main purpose of this utility is to equip local users with tools to monitor backup activities initiated remotely by PPR Administrator. Additionally, it enables to back up contents of the machine to a local disk or network share to later restore individual files or entire volumes. If the target machine runs Microsoft

Hyper-V and has the Hyper-V Application plug-in installed, the Tray Application utility can also be used to agentlessly protect virtual machines hosted by this hypervisor.

Using PPR Installer

1. The product comes in one .exe file. Click on it to initiate the installation.
2. First the installation wizard checks the machine for the presence of Microsoft .Net Framework 4.0 and prompts to install it if required.

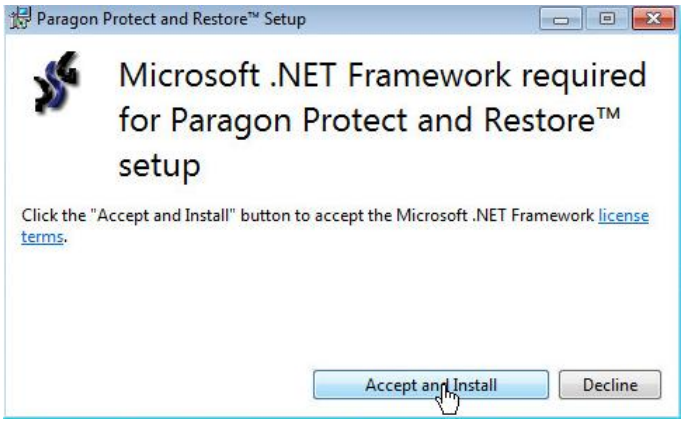

The setup language depends on the system local settings. Currently PPR supports English, German and Russian languages.

3. In order to continue the installation, please read and accept all conditions of Paragon's license agreement by selecting the appropriate option.

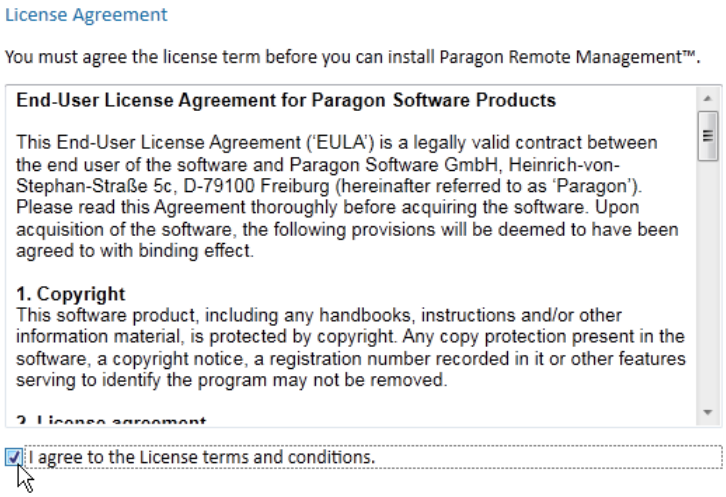

4. Choose the preferred installation mode from two available options:

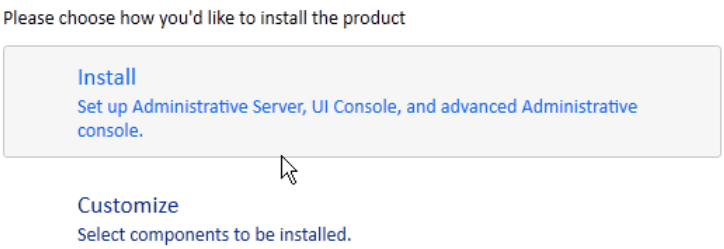

- **Install**. Use this mode to deploy mandatory PPR components in order to start working with the infrastructure. These include PPR Server (the main service that manages all other members of the infrastructure and controls usage of shared resources) and PPR Console (actually two consoles, one GUI and one command-line based on PowerShell). Once you've got these two components installed, all others can be deployed centrally through one of the consoles.
- **Customize**. Use this mode to deploy a particular PPR component(s) from the list of available for your license. This option can save time if you need to install on one machine not only PPR Server and PPR Console, but other components. Besides this option can help if there's no way to do the remote installation.

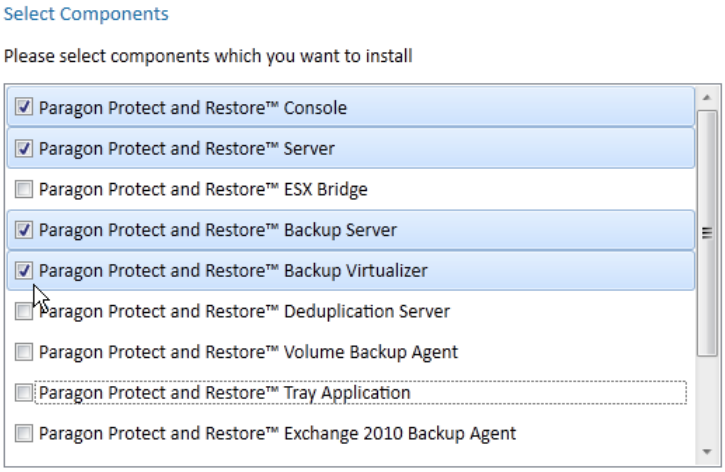

If you don't install PPR Server at this stage, but any other PPR component (except for PPR Console and PPR Installation Client), the wizard will prompt you to provide a DNS name or IP address of a machine where PPR Server has been installed and credentials to access it. If this computer should be accessed with some other credentials, please also provide them after marking a corresponding checkbox.

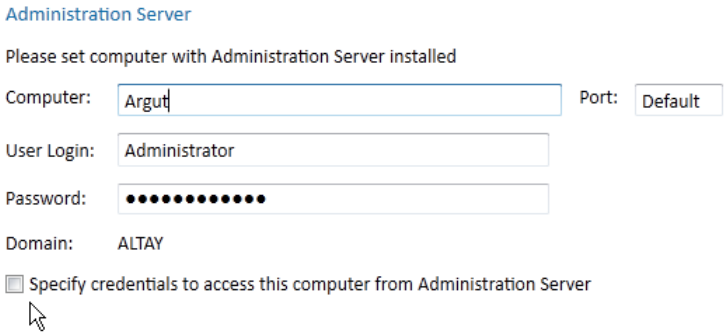

5. Next you need to select whether MS Active Directory facilities (**In Domain**) or credentials of a local machine (**In Workgroup**) will be used to authenticate users to grant access to the PPR infrastructure.

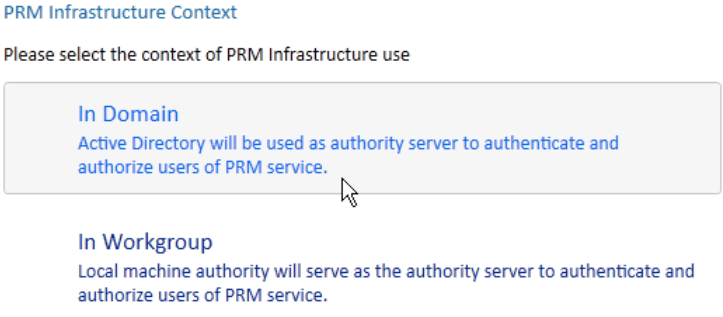

6. Depending on your choice at the previous step, provide credentials of a domain or a local administrator. Please make sure the domain administrator also joins the 'local admins' group.

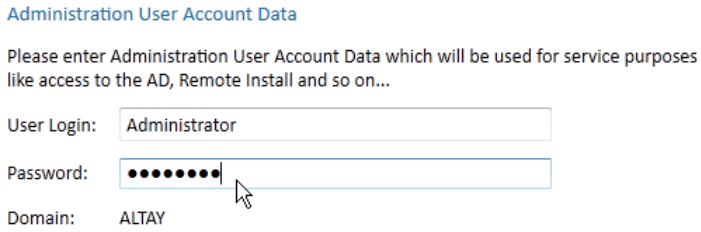

Beside domain or local administrators, PPR can also be administered by members of special groups created during the installation. For more information, please click here.

7. If selecting Active Directory environment (**In Domain** option), the setup wizard will display all of its organizational units (OU) prompting you to choose where PPR's organizational unit with the default security groups should be created.

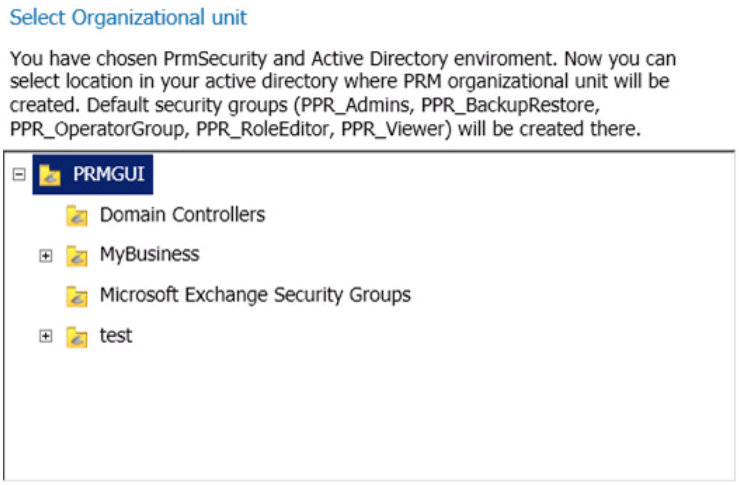

8. The setup wizard will accomplish installation of PPR components according to the specified parameters, which may take several minutes. The wizard will also analyze the computer for required prerequisites and additionally install them if necessary.

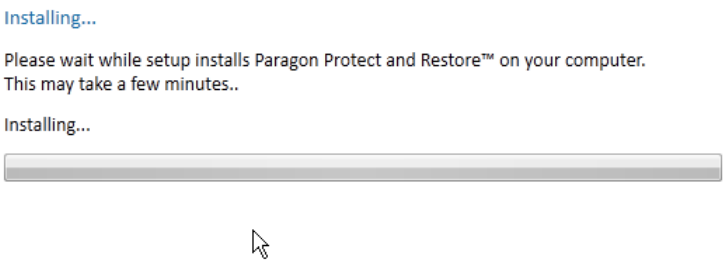

Installing PPR Server and PPR Console

It's recommended to deploy PPR in a Windows domain environment as this configuration offers more options and flexibility.

1. Choose a computer to install **Protect & Restore Server** on. It can be any domain or workgroup machine, but the more powerful, the better.
2. Make sure it meets the systems requirements.
3. Run PPR Installer and go through the default installation procedure.
4. The main indicator that PPR Server has been successfully installed is the **PRM service** running – you can check it up in **Windows Task Manager**.

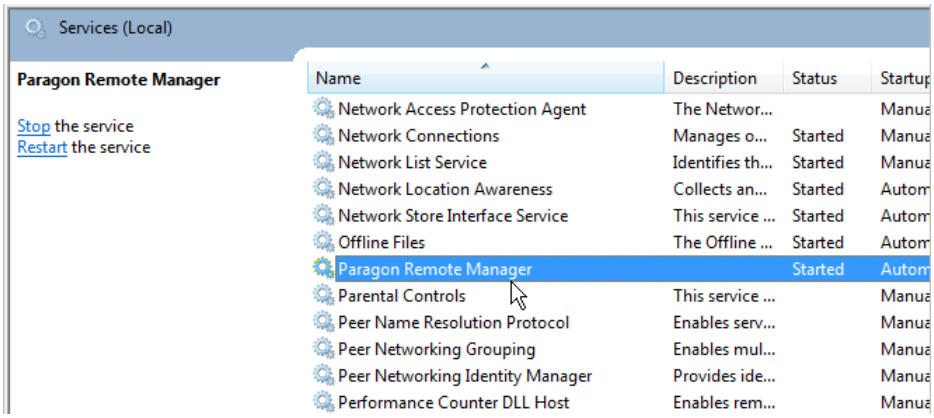

If you're planning to deploy PPR Server and PPR Console on different machines, use a custom mode of the PPR Installer.

Launching Console

1. Launch **Protect & Restore Console**, by either clicking on its desktop icon or going to **Start > Programs > Protect & Restore Console**.
2. Provide a DNS name or IP address of a machine where PPR Server has been installed and access credentials. Click **Connect**.

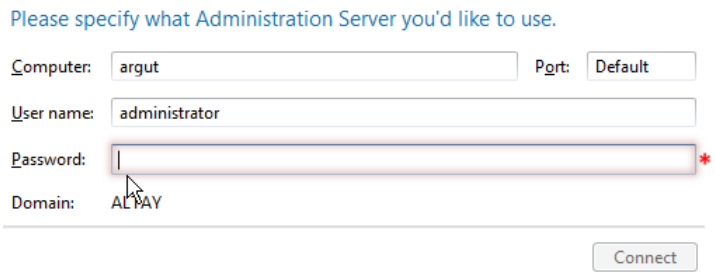

Installing ESX Agent

To get the maximum backup or replication performance, the ESX agent that interacts with the VMware infrastructure should be located as close to datastores containing protected machines as possible. Obviously, when having to do with several vSpheres or ESX hosts, this requirement is not possible to comply via one entry point. In PPR this problem is addressed through an intelligent mechanism, which lets to deploy several ESX agents, one for each datastore. The system then automatically chooses one of the agents as primary to delegate it the management role.

Prerequisites

- Windows Automount should be disabled on the target machine. This requirement is relevant for all versions of Windows OS since 2003. You can disable Windows Automount by launching Windows CMD as Administrator, then entering '**diskpart > automount disable**'.

You can install ESX Agent either centrally through Protect & Restore Console (recommended) or manually on-site. Obviously the manual option is not default, and should only be used if there's no way to do the remote installation.

Through Console

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon (active by default), then select **Add ESX Agents**, or go to **Infrastructure > ESX Agents**, then click on **Configure ESX agent now**.

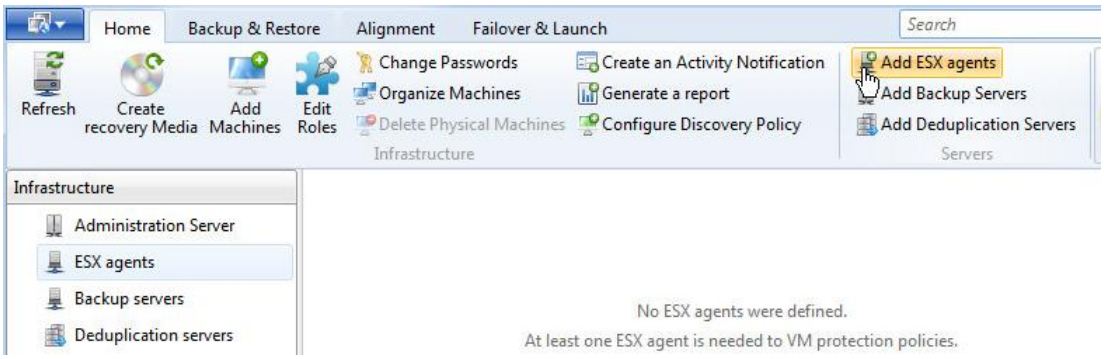

3. Choose a machine to install ESX Agent on. Despite the fact that it can be any machine, the maximum performance and stability can be achieved only when ESX Agent is installed in a guest environment of an ESX server that hosts virtual machines you're planning to protect.

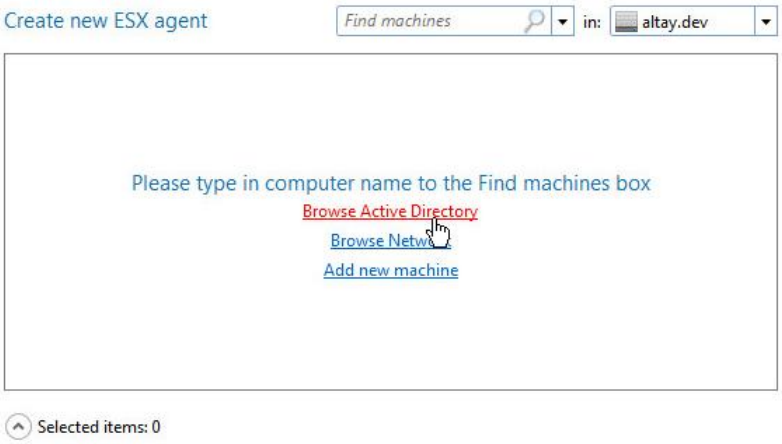

Through the AD browser:

- Click the **Browse Active Directory** hyperlink.
- By default, there only listed machines registered in Active Directory that join a forest domain where PPR is deployed. Select the required machine, then click **Finish** to confirm.

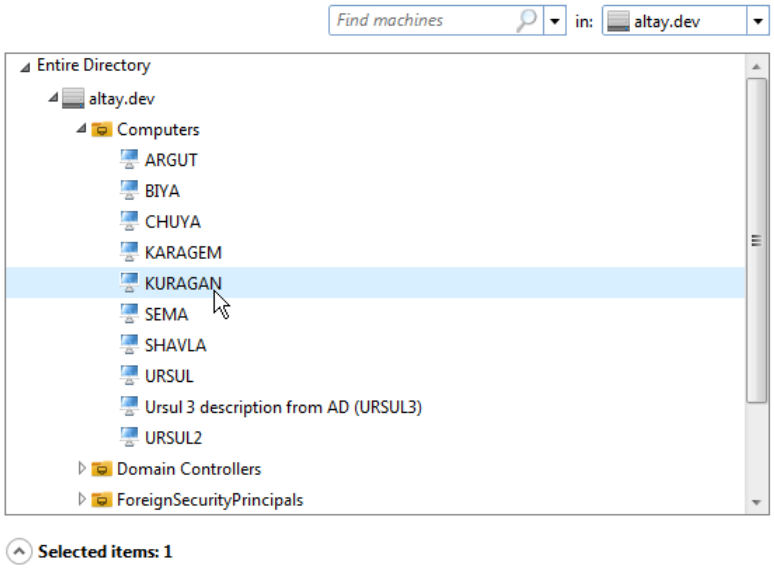

If you'd like to specify a machine that belongs to another forest click on the arrow button at the top right corner of the window, then the **Change forest** link.

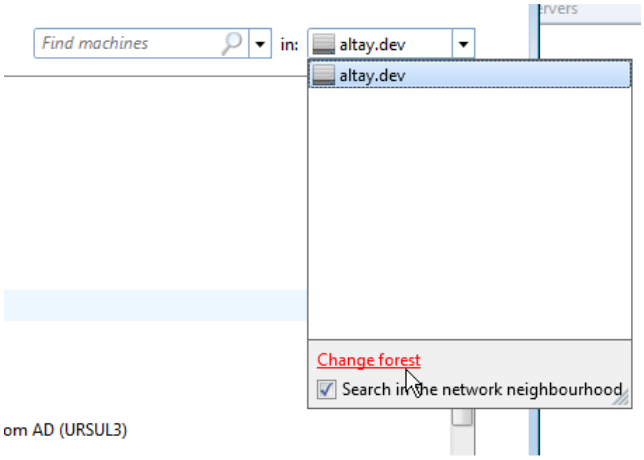

In the opened dialog enter the required forest name**.** Click **Change** when ready.

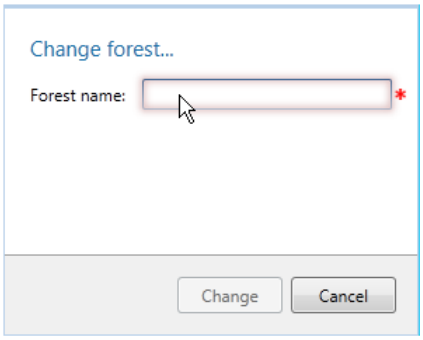

To specify a machine that belongs to some workgroup either use the network browser (convenient, but may take some time to enumerate all machines on the net), or manually provide a name or IP address of the required machine and its access credentials (recommended).

Through the Network browser:

- Click the **Browse Network** hyperlink.
- Wait till all machines on the net are browsed, then select the required one from the list. Please use the **Find machines** field to quickly find the required machine.

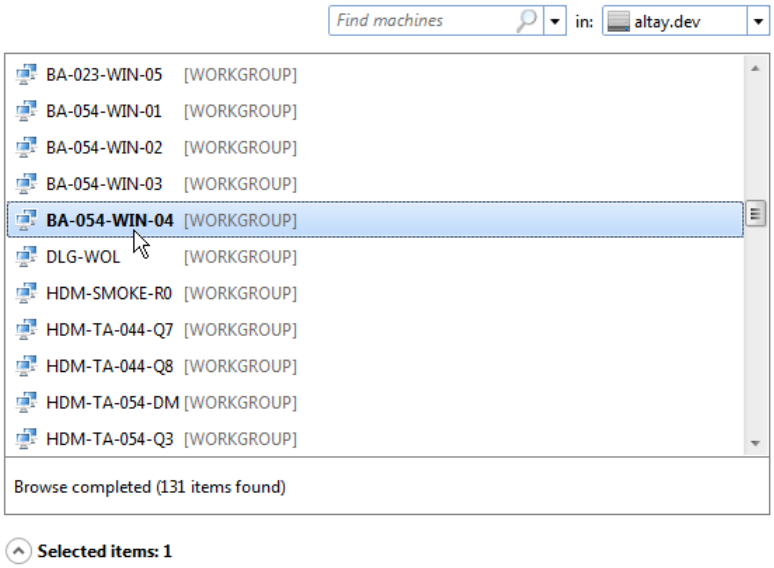

- Provide access credentials of a local administrator. Click **Finish** when ready.

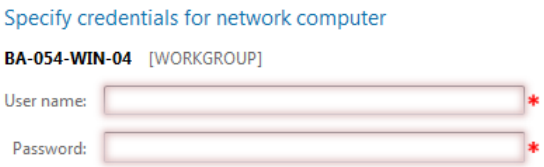

Through the Add new machine dialog:

- Click the **Add new machine** hyperlink.
- Provide a name or IP address of the required machine and its access credentials. Click **OK** when ready.

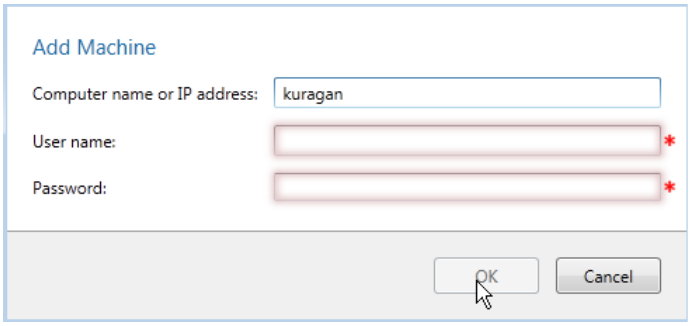

4. There will be created a summary page and a new installation policy, submitted immediately, which you can see through a popup window.

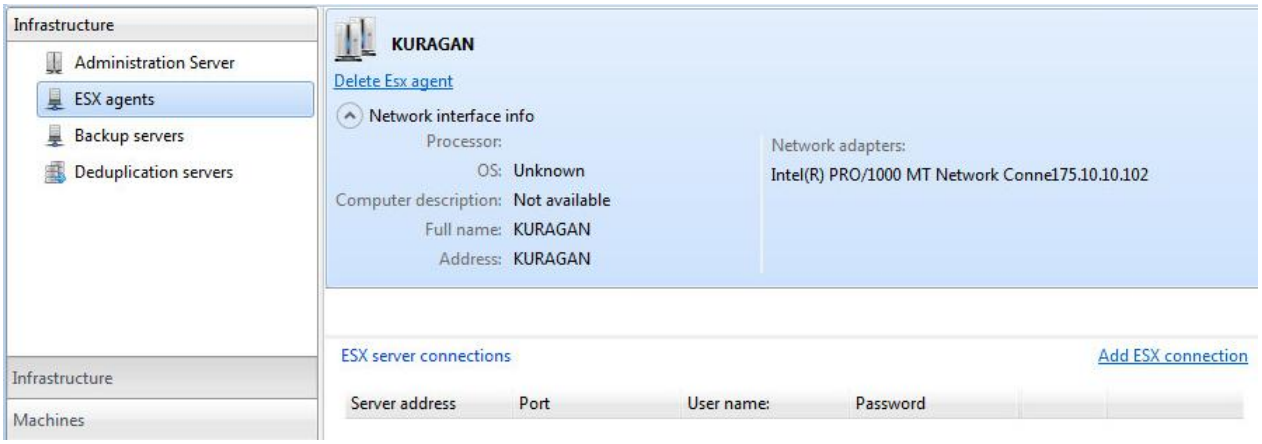

5. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
6. The required machine will be added to the list of computers that join the infrastructure and eventually acquire the roles of **Agent/ESX Agent**.

To make ESX Agent access the VMware infrastructure, you should add at least one ESX host or vSphere. Please consult the Managing ESX Agent chapter for more information.

Manually

1. Choose a machine to install **ESX Agent** on. Despite the fact that it can be any machine, the maximum performance and stability can be achieved only when ESX Agent is installed in a guest environment of an ESX server that hosts virtual machines you're planning to protect.
2. Make sure it meets the systems requirements.
3. Run PPR Installer and go through a custom mode of the PPR Installer.
4. The main indicator that ESX Agent has been successfully installed is appearance of the computer in Protect & Restore Console in the online state. To check it out launch the console, then select **Machines > Physical Machines** to see the required computer.

To make ESX Agent access the VMware infrastructure, you should add at least one ESX host or vSphere. Please consult the Managing ESX Agent chapter for more information.

Installing Backup Server

You can install **Backup Server** either centrally through Protect & Restore Console (recommended) or manually on-site. Obviously the manual option is not default, and should only be used if there's no way to do the remote installation.

Through Console

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon (active by default), then select **Add Backup Servers**, or go to **Infrastructure > Backup Servers**, then click on **Set up Backup Server now**.

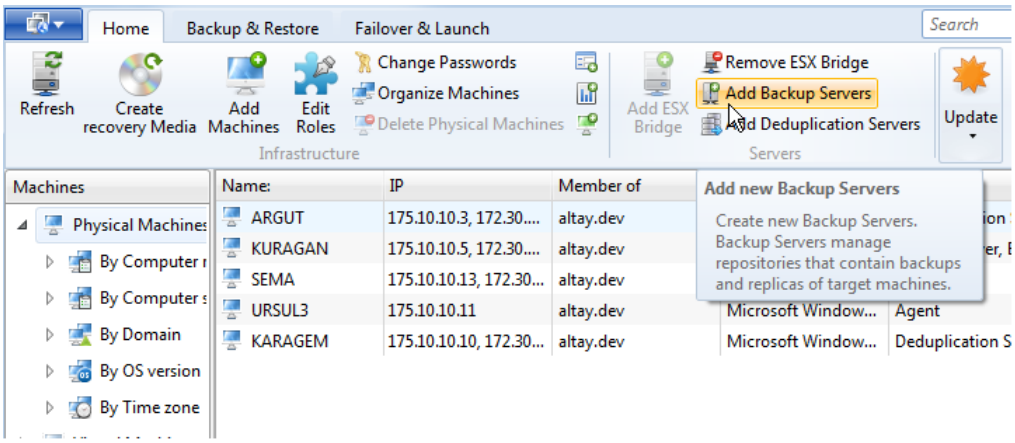

3. Choose a domain or a workgroup machine you're planning to use as Backup Server. The procedure is similar to the remote installation of ESX Agent, so please consult the corresponding scenario for more information.

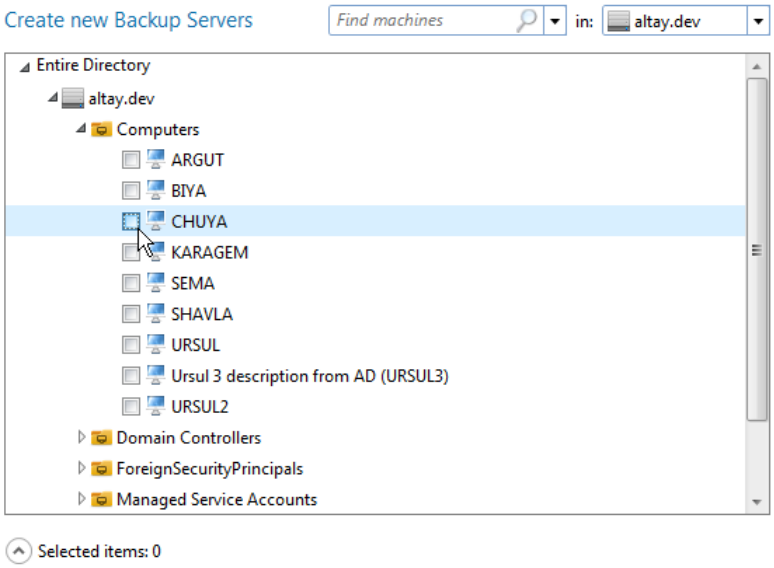

4. There will be created a summary page and a new installation policy, submitted immediately, which you can see through a popup window.
5. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
6. The required machine will be added to the list of computers that join the infrastructure and eventually acquire the roles of **Agent/Backup Server/ESX Storage plug-in/Remote Storage plug-in**.

If you'd like to install Backup Server on an existing member of the infrastructure, please use the Edit Roles dialog. As a result of this operation the PPR service will be restarted.

If you'd like to launch machines directly out of backup images in VMware ESX environment, please use the Edit Roles dialog to additionally install the Backup Virtualizer plug-in.

To allow Backup Server to store backups of physical and virtual machines or VM replicas, you should configure corresponding primary storages, a local or network one for backups, and an ESX one for replicas. Please consult the Configuring Backup and Replica Storages chapter for more information.

Backup Server can also be deployed through the Add Machines Wizard.

Manually

The procedure is similar to the manual installation of ESX Agent, so please consult the corresponding scenario for more information.

Installing Deduplication Server

You can install **Deduplication Server** either centrally through Protect & Restore Console (recommended) or manually on-site. Obviously the manual option is not default, and should only be used if there's no way to do the remote installation.

Through Console

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon (active by default), then select **Add Deduplication Servers**, or go to **Infrastructure > Deduplication Servers**, then click on **Add Deduplication Servers**.

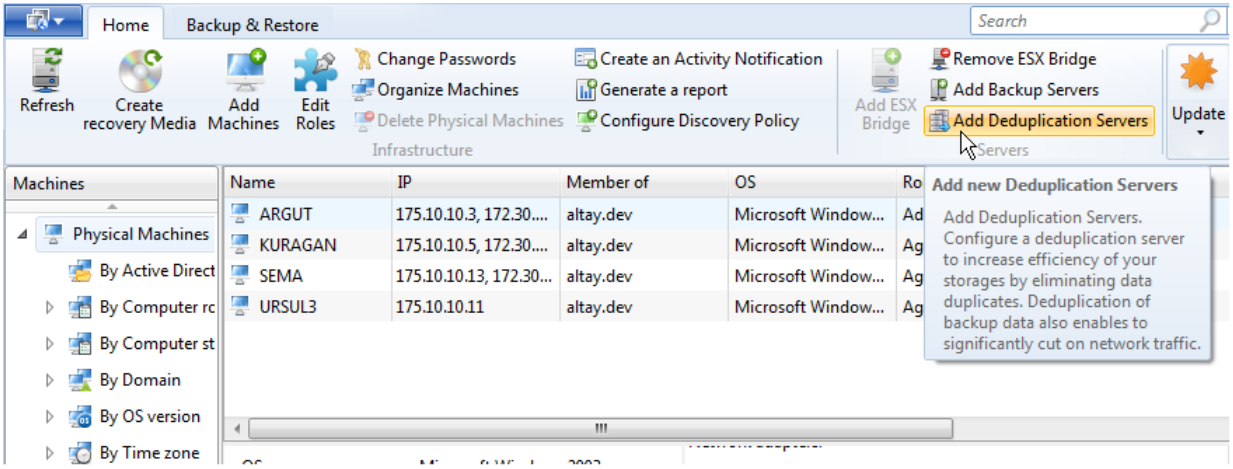

3. Choose a domain or a workgroup machine you're planning to use as Deduplication Server. The procedure is similar to the remote installation of ESX Agent, so please consult the corresponding scenario for more information.

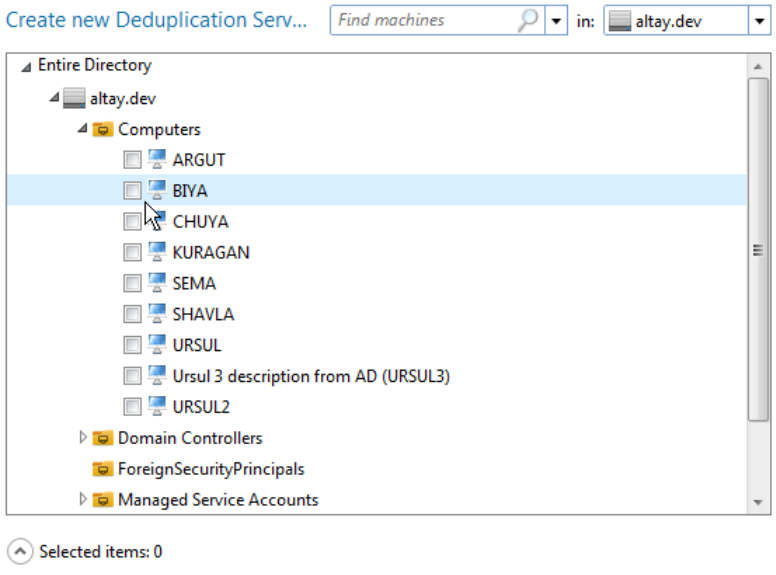

4. There will be created a summary page and a new installation policy, submitted immediately, which you can see through a popup window.
5. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
6. The required machine will be added to the list of computers that join the infrastructure and eventually acquire the roles of **Agent/Deduplication Server**.

If you'd like to install Deduplication Server on an existing member of the infrastructure, please, please use the Edit Roles dialog. As a result of this operation the PPR service will be restarted.

To let Deduplication Server work, you need to configure a deduplication storage and link existing backup storages to the server.

Deduplication Server can also be deployed through the Add Machines Wizard.

Manually

The procedure is similar to the manual installation of ESX Agent, so please consult the corresponding scenario for more information.

Adding Target Machines

To do agent based protection of Windows machines (physical or virtual), you need to add them to the infrastructure. It can be done either centrally through Protect & Restore Console (recommended) or manually by installing **Backup Agent** on all target computers. Obviously the manual option is not default, and should only be used if there's no way to do the remote installation.

If you'd like to do agent based protection of a machine that already has the role of Administration Server, Backup Server, or both, please use the Edit Roles dialog, as machines that already join the infrastructure with one of the mentioned above roles will not be available in the Add Machines Wizard. As another option, you can always install Backup Agent manually on-site.

Through Console

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon then select **Add Machines**.

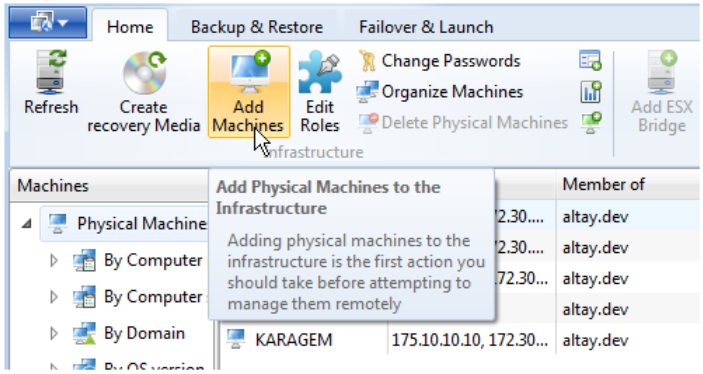

3. Choose domain or workgroup machines you're planning to add to the PPR infrastructure. The procedure is similar to the remote installation of ESX Agent, so please consult the corresponding scenario for more information.

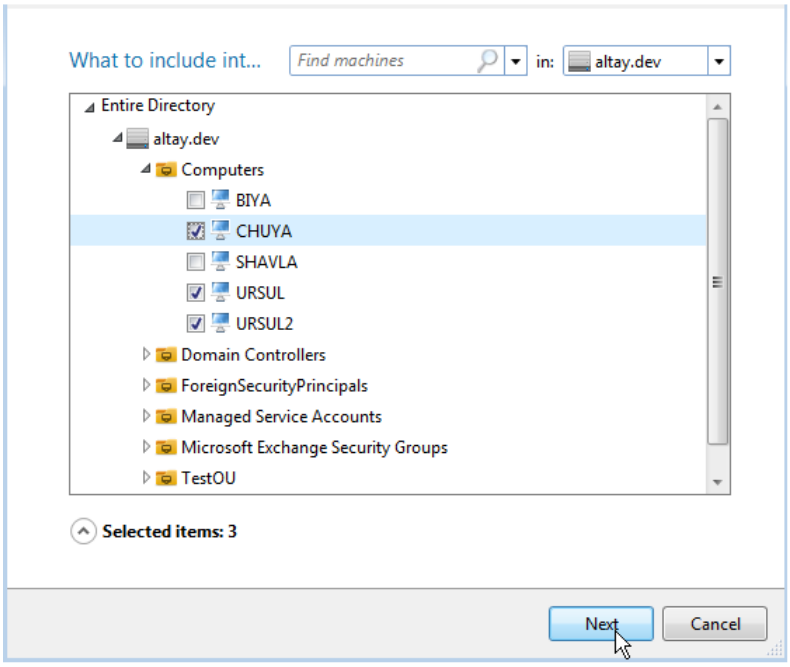

The wizard lists both physical and virtual machines. You don't need to add guests of VMware ESX to the PPR infrastructure to protect them, as this can be done agentlessly through ESX Agent (much higher performance, complete data safety, minimal downtime, etc.). However, if you're dealing with a VMware fault-tolerant system that cannot be protected agentlessly, you can add this machine to the infrastructure to protect it through the on-site backup agent.

4. Specify roles to install. By default, the wizard will only add all selected machines to the infrastructure, assigning them the general role of Agent. To allow their protection, you should extend functionality of this agent by selecting required plug-in(s). The number of available plug-ins depends on the purchased license. Let's take a detailed look at all possible plug-ins:

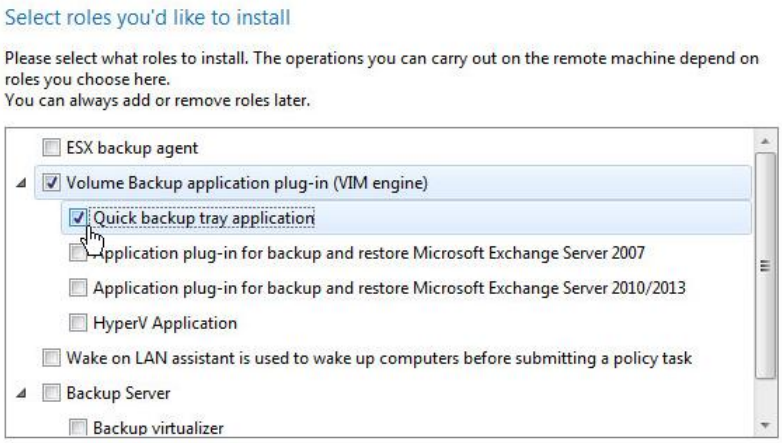

- **ESX Backup Agent**. It embeds the ESX agent role to all selected machines.
- **Volume backup application plug-in**. It extends functionality of the general agent allowing it to back up an entire machine or separate volumes at sector level. It also enables to do remote restore of data volumes.
  - **Exchange Server 2007 application plug-in**. It extends functionality of the general agent allowing it to back up and restore databases of MS Exchange Server 2007.
  - **Exchange Server 2010/2013 application plug-in**. It extends functionality of the general agent allowing it to back up and restore databases of MS Exchange Server 2010/2013.
  - **Hyper-V Application**. It extends functionality of the general agent allowing it to back up and restore guest machines hosted by Hyper-V.
  - **Quick backup tray application**. It allows the user to create backup images on his own to later restore individual files or entire volumes. Besides through this utility the user can monitor backup activities initiated remotely on his computer by PPR Administrator.
- **Wake-on-LAN assistant**. It doesn't provide any backup functions. It only allows waking up remote target machines to do backup. There's no need to install this plug-in on all machines that require wake-up-on-LAN. Just select one that is always online and it will automatically wake up all others that share the same subnet when needed.
- **Backup Server**. It embeds the backup server role to all selected machines.
  - **Backup virtualizer**. It extends functionality of the backup server, allowing it to launch machines directly out of backup images in the VMware ESX environment through the NFS (Network File System) server facilities.
- **Deduplication Server**. It embeds the deduplication server role to all selected machines.

The number of available plug-ins depends on the purchased license.

5. Specify how and when the selected plug-ins should be installed. You've got three options to choose from:

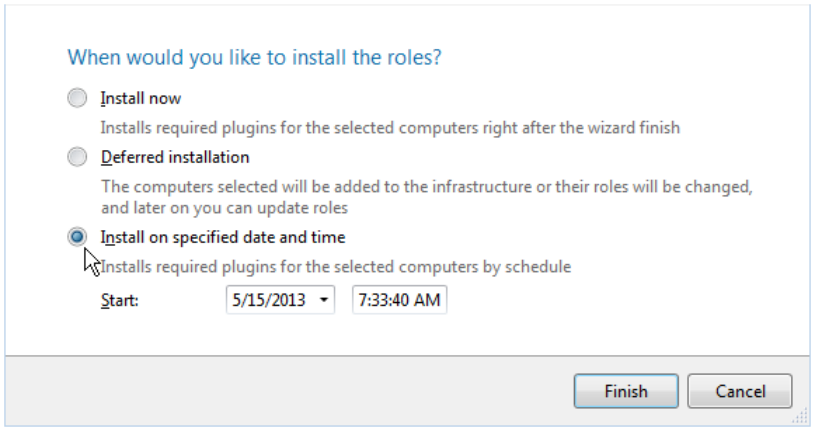

- **Install now**. Installation will start immediately after confirmation.
- **Deferred installation**. If you choose this option all target machines will become available in the infrastructure with the statuses 'Role pending', 'Off-line' as show below.

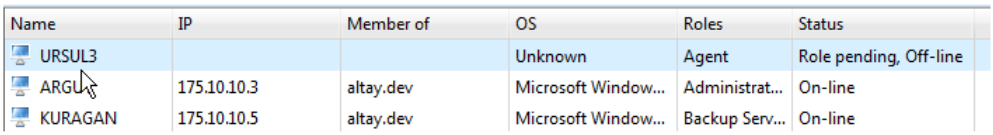

Name	IP	Member of	OS	Roles	Status
URSUL3			Unknown	Agent	Role pending, Off-line
ARGUT	175.10.10.3	altay.dev	Microsoft Window...	Administrat...	On-line
KURAGAN	175.10.10.5	altay.dev	Microsoft Window...	Backup Serv...	On-line

That means no installation procedures will be initiated until you right click each machine and then select **Update roles**.

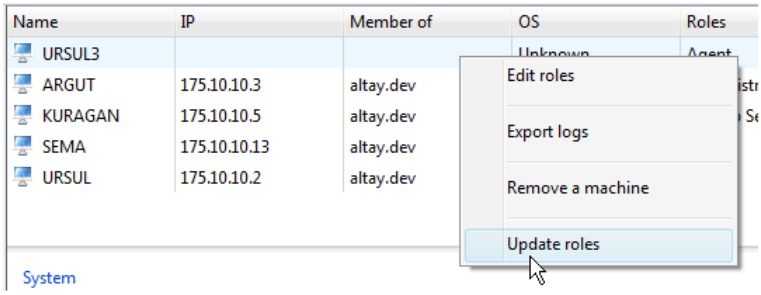

- **Install on specified date and time**. By choosing this option you can schedule installation of roles on target machines.

6. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
7. The required machine(s) will be added to the list of computers that join the infrastructure (select **Machines > Physical Machines**) and eventually acquire the specified roles.

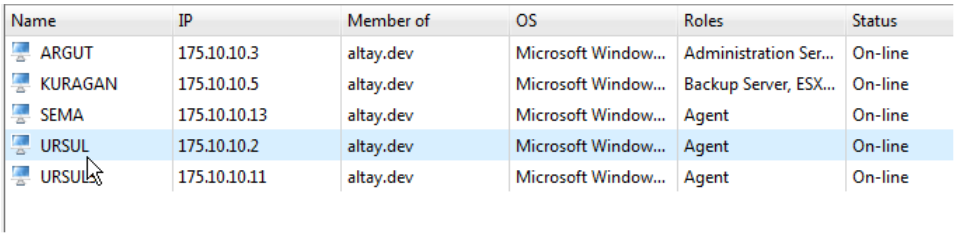

Name	IP	Member of	OS	Roles	Status
ARGUT	175.10.10.3	altay.dev	Microsoft Window...	Administration Ser...	On-line
KURAGAN	175.10.10.5	altay.dev	Microsoft Window...	Backup Server, ESX...	On-line
SEMA	175.10.10.13	altay.dev	Microsoft Window...	Agent	On-line
URSUL	175.10.10.2	altay.dev	Microsoft Window...	Agent	On-line
URSUL	175.10.10.11	altay.dev	Microsoft Window...	Agent	On-line

Automatically through policy

1. Launch Protect & Restore Console.

2. If a connection with the server has been established, click on the **Home** ribbon then select **Configure Discovery Policy**.

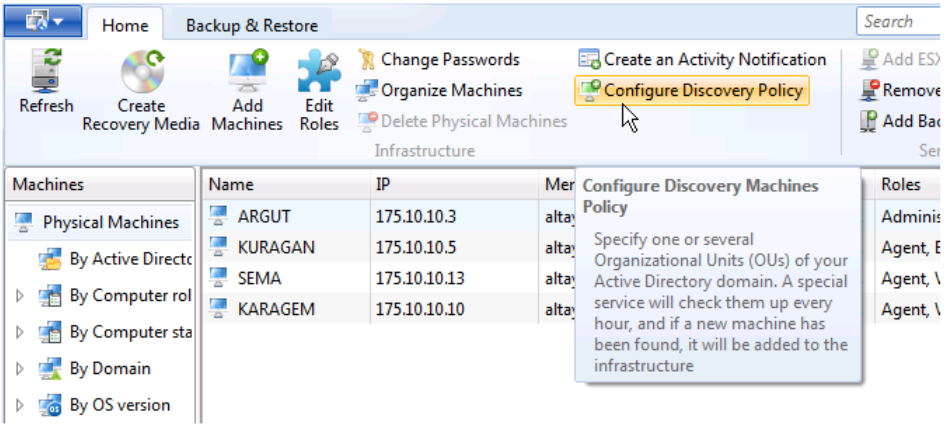

3. Either manually enter a name of the required organizational unit or AD container in the 'Search' field or click on the **Browse Active Directory** hyperlink to navigate through all OUs of your Active Directory.

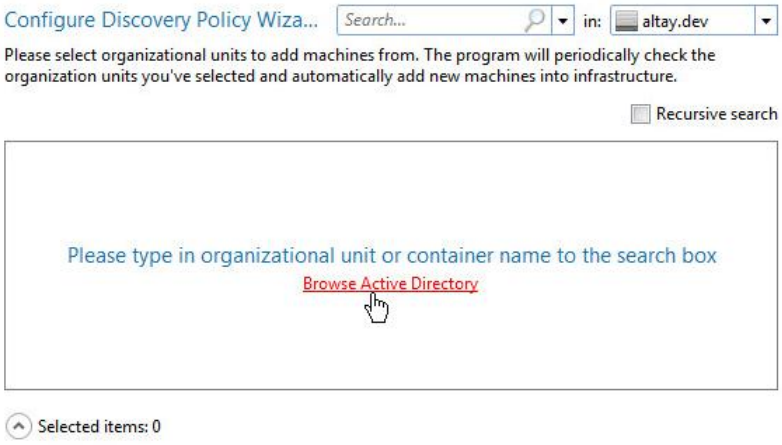

If you've got too many organizational units in AD, their browsing and listing may take plenty of time. Please use the 'Search' field if you know exactly what OUs or containers you'd like to search for.

4. Specify one or several units or containers you'd like our program to check for new computers. Mark the **Recursive search** option if you'd like the program to process not only root but all subfolders inside selected objects.

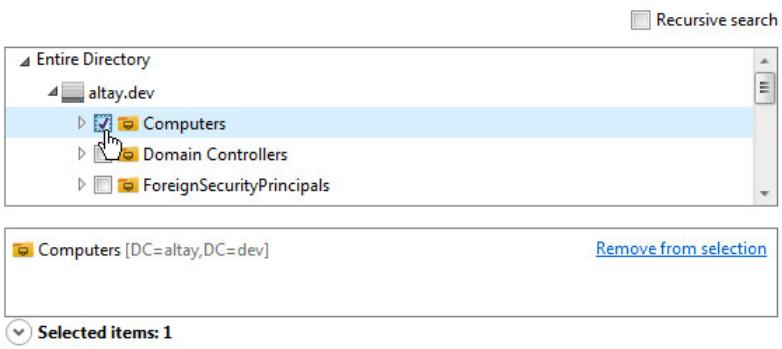

5. Choose how often you'd like a special service to check the selected objects for new machines (every hour by default) and a number of additional parameters.

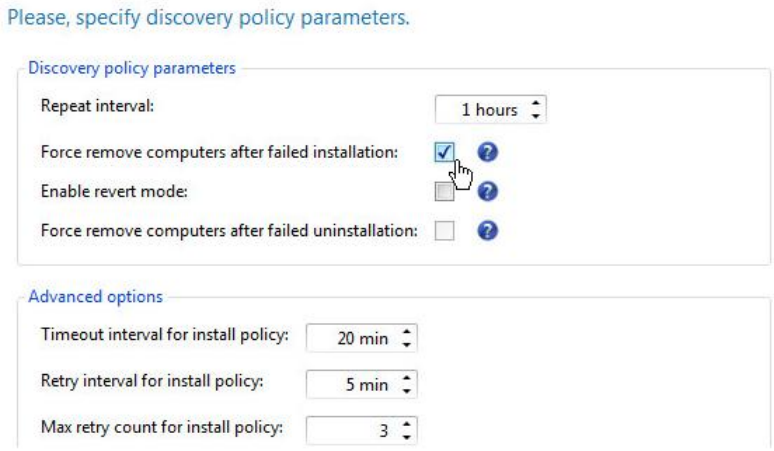

Discovery policy parameters

- **Repeat interval**. By default selected objects will be checked for new machines every hour, which you can change according to your needs.
- **Force remove computers after failed installation**. If a machine hasn't been added to the PPR infrastructure (Firewall restrictions, incorrect credentials, some compatibility issues, etc.), it still remains in the list of the infrastructure members (**Infrastructure > Machines**). Please mark the option to automatically clear this list from invalid members.
- **Enable revert mode**. If this mode is active, members that do not join specified organizational units or AD containers will be automatically removed from the infrastructure.
- **Force remove computers after failed uninstallation**. If a machine hasn't been removed from the PPR infrastructure with a success, it still remains in the list of the infrastructure members (**Infrastructure > Machines**). Please mark the option to automatically remove this type of machines from the infrastructure.

Advanced options

- **Timeout interval…** By default, there are 20 minutes for a discovery installation policy to be completed on a target machine. Change the default value if necessary. Please note if you set this value to zero and one or several target machines are offline or temporarily unavailable, the installation policy will fail immediately with a corresponding error.
- **Retry interval…** By default, a failed discovery installation policy will be re-submitted in five minutes if two or more attempts are specified in the next option. Change the default value if necessary.
- **Max retry count…** By default, a failed discovery installation policy will be re-submitted three times pausing between the attempts as specified in the previous option. If no success, it will be aborted with a corresponding error. Change the default value if necessary.

Advanced settings will be available if the corresponding option is enabled in the Settings dialog.

6. Specify plug-ins to install. Click **Finish** when ready.
7. Now if a new machine has been found in the specified organizational unit(s), it will be automatically added to the infrastructure with the required roles.

Manually

The procedure is similar to the manual installation of ESX Agent, so please consult the corresponding scenario for more information.

Installing Tray Application

Tray Application can be installed either centrally through Protect & Restore Console (recommended) or manually on-site. Obviously the manual option is not default, and should only be used if there's no way to do the remote installation.

Through Console

If the required machine has already been added to the infrastructure, then use the **Edit roles** dialog to additionally install the utility on it. If not, install it when adding this machine to the infrastructure.

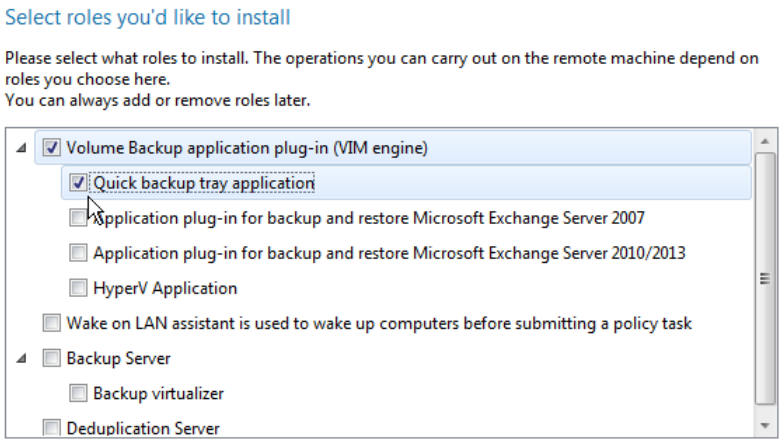

Manually

The procedure is similar to the manual installation of ESX Agent, so please consult the corresponding scenario for more information.

Tray Application requires a Backup Agent role present on the target machine, thus it will be added automatically if necessary.

For more information about Tray Application, please consult Using Tray Application.

Building WinPE Recovery Media

Prerequisites

- A computer where you're going to prepare WinPE recovery media should run Windows Vista x64 or later operating system.
- Before you launch Recovery Media Builder please make sure you've got Windows Assessment and Deployment Kit (ADK) 8.0 or later installed in the chosen system. Otherwise, you won't be able to accomplish the operation. ADK is a Microsoft's proprietary tool and can be obtained from its Download Center for free. Please note that you need a genuine Windows installation to be able to download this tool. Moreover you will need to download a version, which is suitable for your Windows OS – Recovery Media Builder automatically detects your system and offers the required download link.

Preparing the WinPE recovery environment on a flash stick

1. Plug in a thumb drive of at least 250 MBs in size. Please note all data on that drive will be deleted.

2. Launch Protect & Restore Console.
3. If a connection with the server has been established, click on the **Home** ribbon then select **Create recovery media**.

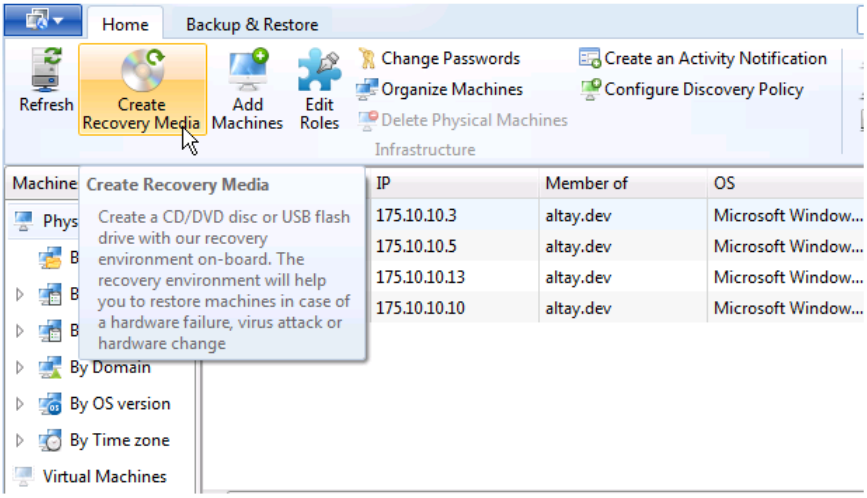

4. First you will be asked to specify Administration Server to connect to during recovery operations.

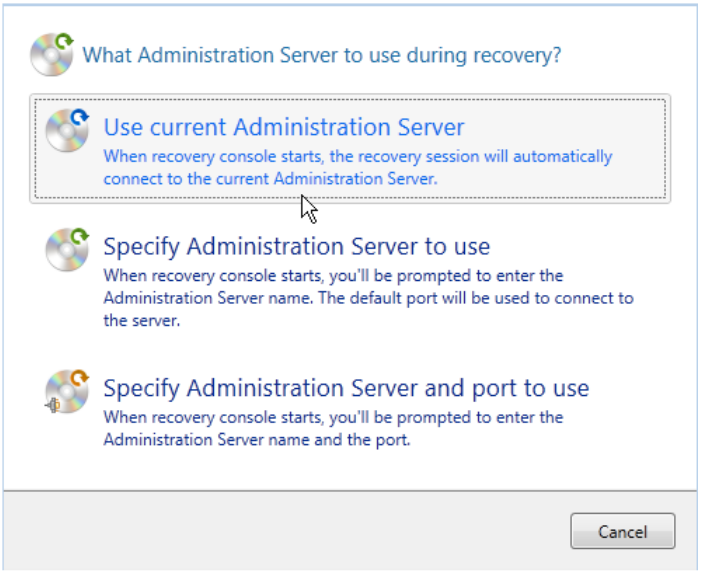

- **Use current Administration Server**. Once the recovery environment has been loaded, there will automatically be established connection to Administration Server used by the console at the moment.
- **Specify Administration Server to use**. Once the recovery environment has been loaded, you will be prompted to specify a DNS name or IP address of Administration Server, you'd like to connect to.
- **Specify Administration Server and port to use**. Similar option to the previous one, but you will be additionally prompted to specify a port, if it's not default.

5. The welcome page introduces the wizard's functionality. Click **Next** to proceed.
6. Click on **Removable flash media**, then select a thumb drive from the list of flash memory devices available in the system at the moment (if several). If you'd like to create an ISO image of the WinPE environment, please use the corresponding option.

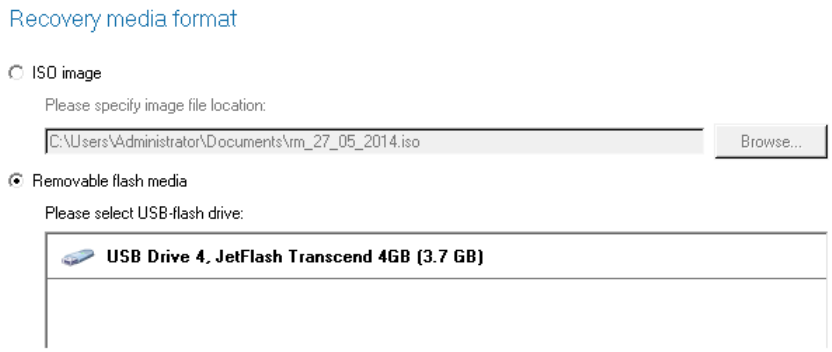

7. The wizard will warn you that all data on the selected drive will be deleted. Please confirm the operation to proceed.

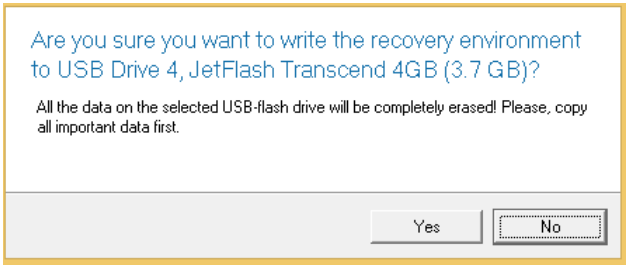

8. Inject additional drivers for specific storage controllers, network cards, or other devices. First click on the required link.

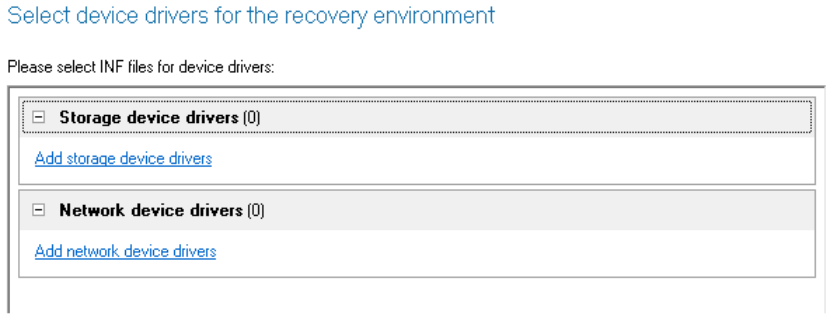

In the opened dialog browse for an .INF file of the required driver package located on a floppy disk, local disk, CD/DVD or a network share.

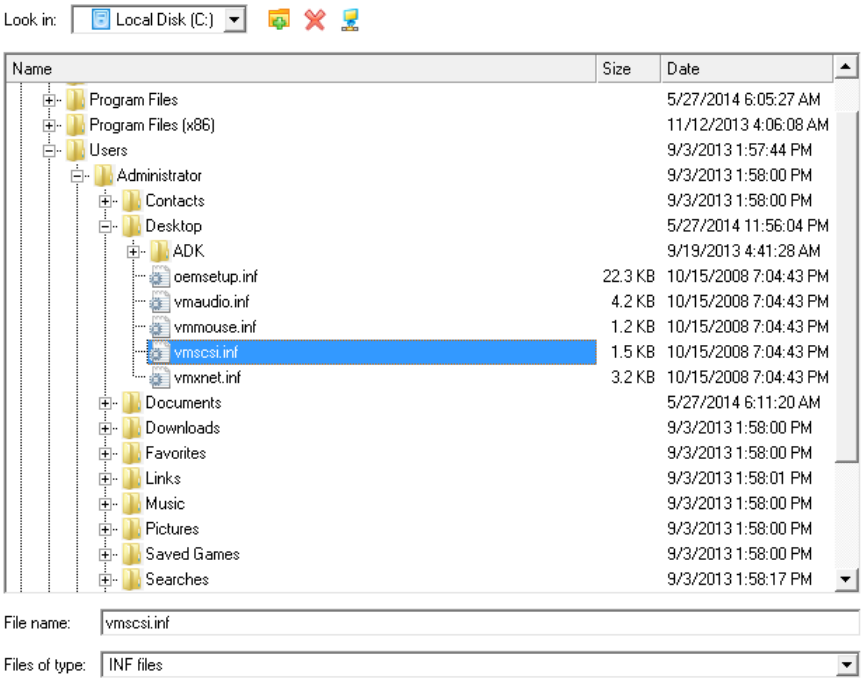

If successfully injected, you can see the specified driver on the list. If you'd like to add another driver, please repeat the procedure.

By default, WinPE environment doesn't support touch screens, thus Windows pads can only be managed by mouse and/or keyboard. You can add required drivers however through our driver injector.

9. Set up a network connection if needed. You've got several options to choose from:

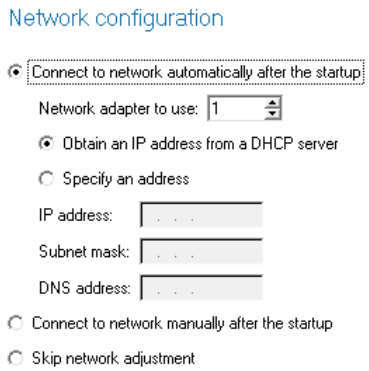

- **Connect to network automatically**... Specify the following parameters if you'd like to have an active network connection once the bootable environment has been started up:
  - **A network adapter to use**. Select a network adapter (if several in the system) to be used for the network connection.
  - **IP address settings.** Choose whether to get an IP address automatically from a DHCP server or set it manually.
- **Connect to network manually**... If selecting this option you will be prompted to configure network properties each time the computer has been started up from the bootable media.
- **Skip network adjustment**. Please use this option if you don't need network support on the bootable media.

10. Map a network share if needed (not be available if selecting **Skip network adjustment** on the previous step). You've got several options to choose from:

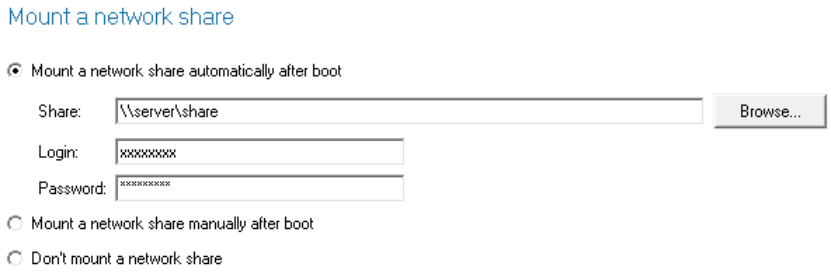

- **Mount a network share automatically**... Mark this option if you'd like to have a mapped network resource once the computer has been started up from the bootable media. Manually type in a path to the required network share or click **Browse** to find it, then provide user credentials if necessary.

With no pre-defined user credentials your network share will be attempted to map under the Guest account.

Please use back slashes for WinPE-based media, like \\server\share.

- **Mount a network share manually**... If selecting this option you will be prompted to map a network resource each time the computer has been started up from the bootable media.
- **Don't mount a network share**. Please use this option if you don't need to map network resources.

11. If you install WAIK or ADK by their default locations, the wizard automatically detects it. Otherwise, you will need to browse for the required folder. If you haven't installed one of these tools yet, click **Download WAIK/ADK** to get them directly from the Microsoft website.

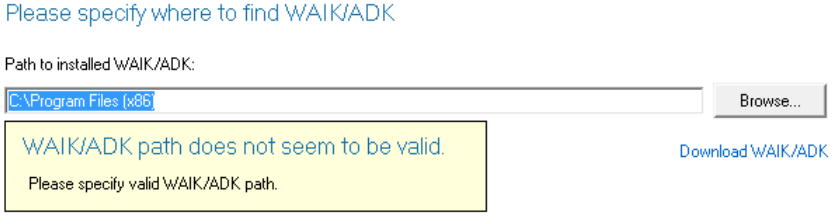

Manually browsing for Windows OS kits:

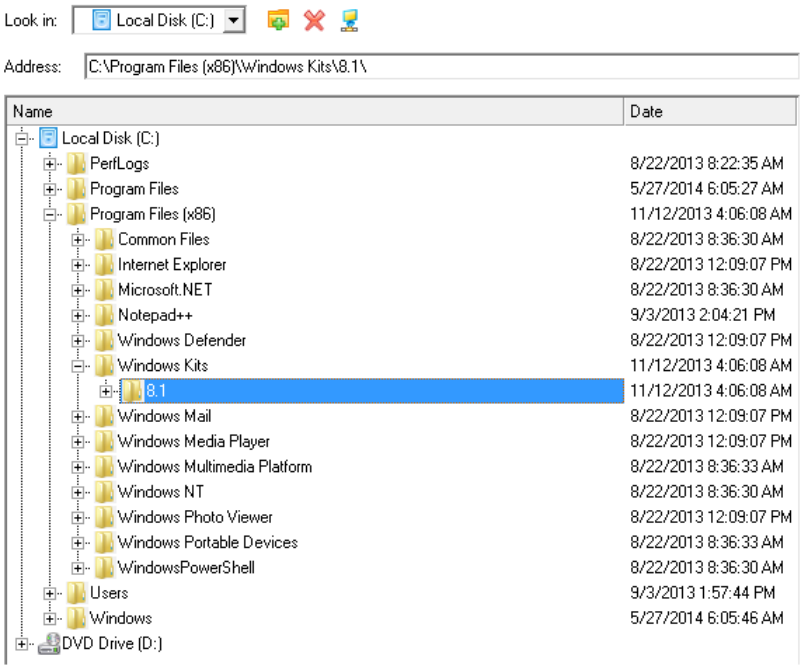

The wizard won't continue until you install WAIK or ADK.

Please take the following information into account:

- **If running RMB under Windows 7, 8, 8.1, Server 2008 R2, Server 2012 R1/R2, please use ADK 8.1;**
- **If running RMB under Windows Vista, Server 2008 R1, please use WAIK of Windows 7**

12. When done with all parameters, click **OK** to initiate creation of the specified bootable media, which takes a couple of minutes.

Typical Use Cases

Managing ESX Agent

Privileges to manage vSphere guests

The security model of VMware allows much flexibility in limiting access and management rights of any object of the virtual infrastructure. In vSphere 5.0 for instance, there are distinguished 255 privileges. To do backup, restore, partition alignment of ESX guests, or to store VM replicas on ESX datastores, PPR may require up to 50 privileges.

Depending on tasks you're going to accomplish with PPR, you can create one or several users in vSphere or modify already existed users according to the information above. For instance, the standard user 'VMBackupUser2' can well be used to do backup of ESX guests. If you don't want to waste time on configuring users with specific privileges, you can use an administrative account of the datacenter you're going to manage + add it the 'Global.Licenses' privilege.

Only one account can be used for one ESX connection, thus if you need to back up and restore ESX guests for instance, you need to provide an account that cover privileges for both operations.

Let's see what privileges are needed for each type of operations:

Privilege	Backup	Align	Store	Restore
Category 'Global'				
Global.CancelTask	+	+	+	+
Global.Licenses	+	+	+	+
Category 'Folder'				
Folder.Create	-	-	+	+
Folder.Delete	-	-	+	+
Category 'Datastore'				
Datastore.Browse	+	+	+	+
Datastore.FileManagement	+	+	+	+
Datastore.AllocateSpace	+	+	+	+
Datastore.UpdateVirtualMachineFiles	-	+	+	+
Category 'Network'				
Network.Assign	-	-	-	+
Category 'Host > Configuration'				
Host.Config.Storage	+	+	+	+

Category 'Virtual machine > Inventory'				
VirtualMachine.Inventory.Create	-	-	+	+
VirtualMachine.Inventory.Delete	-	-	+	+
VirtualMachine.Inventory.Move	-	-	-	+
Category 'Virtual machine > Interaction'				
VirtualMachine.Interact.PowerOn	-	+	-	+
VirtualMachine.Interact.PowerOff	-	+	-	+
VirtualMachine.Interact.DeviceConnection	-	-	+	+
Category 'Virtual machine > Configuration'				
VirtualMachine.Config.Rename	-	-	+	+
VirtualMachine.Config.AddExistingDisk	-	-	+	+
VirtualMachine.Config.AddNewDisk	-	-	+	+
VirtualMachine.Config.RemoveDisk	-	-	+	+
VirtualMachine.Config.CPUCount	-	-	+	+
VirtualMachine.Config.Memory	-	-	+	+
VirtualMachine.Config.AddRemoveDevice	-	-	+	+
VirtualMachine.Config.Settings	+	+	+	+
VirtualMachine.Config.Resource	-	+	+	+
VirtualMachine.Config.DiskLease	+	+	+	+
VirtualMachine.Config.ChangeTracking	+	-	-	-
Category 'Virtual machine > State'				
VirtualMachine.State.CreateSnapshot	+	+	+	+
VirtualMachine.State.RevertToSnapshot	-	+	+	+
VirtualMachine.State.RemoveSnapshot	+	+	+	+
VirtualMachine.State.RenameSnapshot	+	+	+	+
Category 'Virtual machine > Provisioning'				
VirtualMachine.Provisioning.Clone	-	-	-	+
VirtualMachine.Provisioning.DiskRandomAccess	-	+	+	+
VirtualMachine.Provisioning.DiskRandomRead	+	+	+	+
Category 'Resource'				

Resource.AssignVMToPool	-	-	+	+
Resource.CreatePool	-	-	+	+
Resource.RenamePool	-	-	+	+
Resource.EditPool	+	+	+	+
Resource.DeletePool	-	-	+	-
Resource.HotMigrate	-	-	-	+
Resource.ColdMigrate	-	-	-	+

If you'd like to know how to create users with specific privileges in vSphere, please consult documentation provided by VMware.

Adding ESX connections

PPR can simultaneously work with several ESX hosts, or vCenters. To register a standalone ESX host or vCenter, please do the following:

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore ESX Agent** is installed.
- Enough ESX privileges.

1. Once you've got ESX Agent installed, select **Infrastructure > ESX Agents**, then click on **Add ESX connection**.

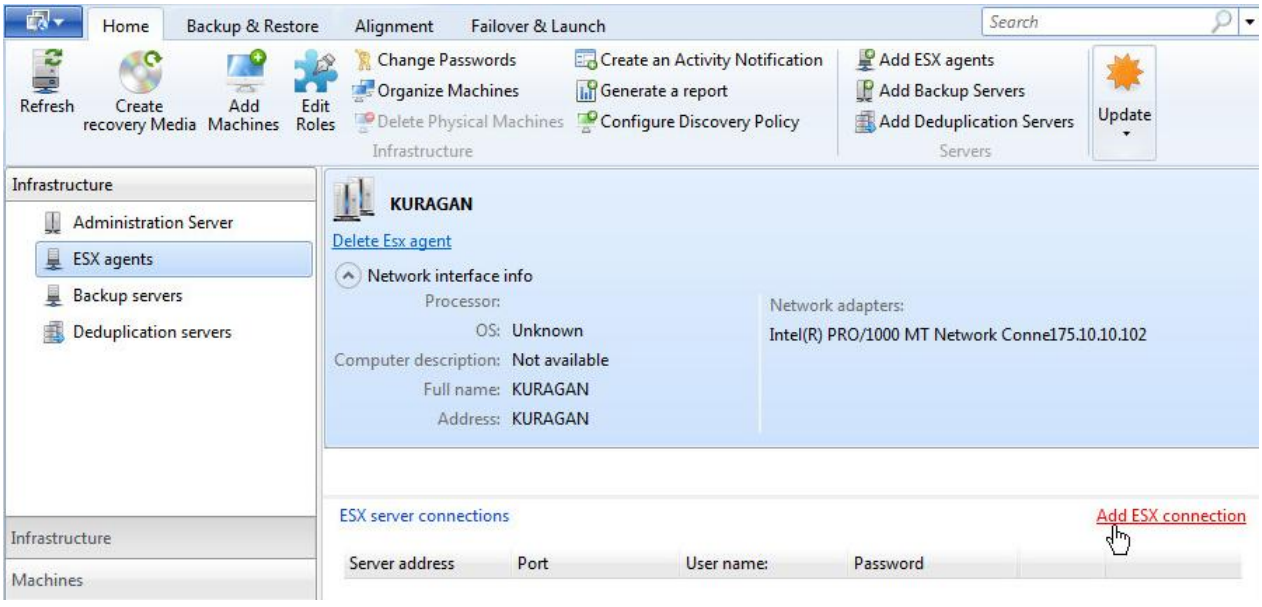

2. Enter a DNS name or IP address of the required vCenter or ESX host, a communication port (if necessary), and administrator credentials in the corresponding fields. Click **Save changes** to start the operation, which takes approx. a couple of seconds. If a success, the corresponding connection will be added to the list.

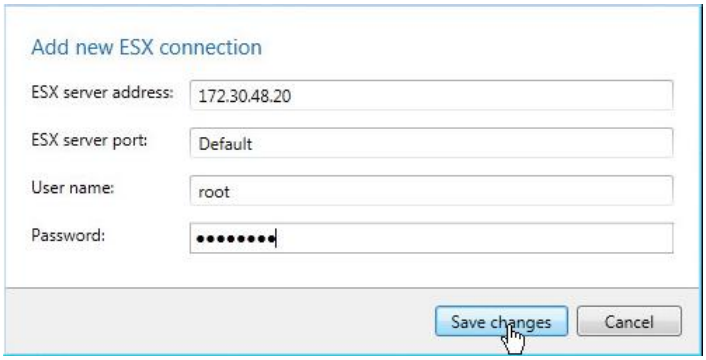

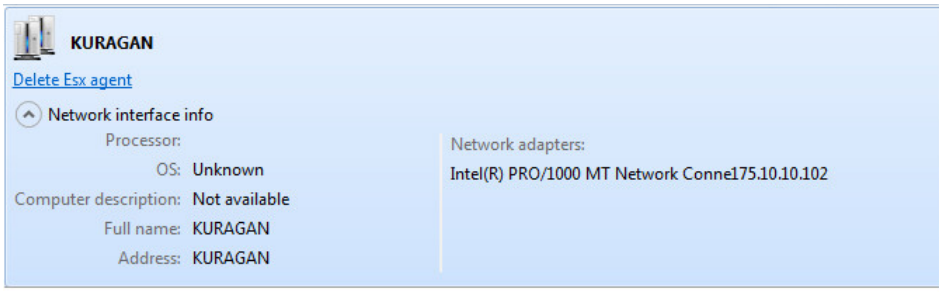

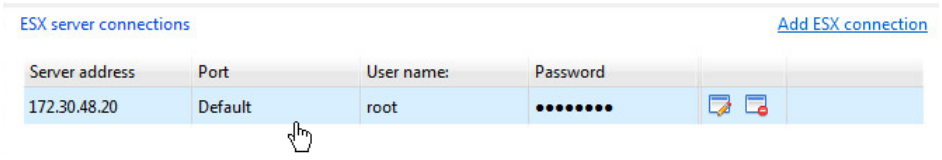

If the required ESX host is a member of a vCenter, always use the IP address and credentials of that vCenter.

3. Select **Machines > Virtual Machines > Host and Clusters** to see the connected ESX host, or vCenter. It's the main indicator that ESX Agent has been appropriately configured.

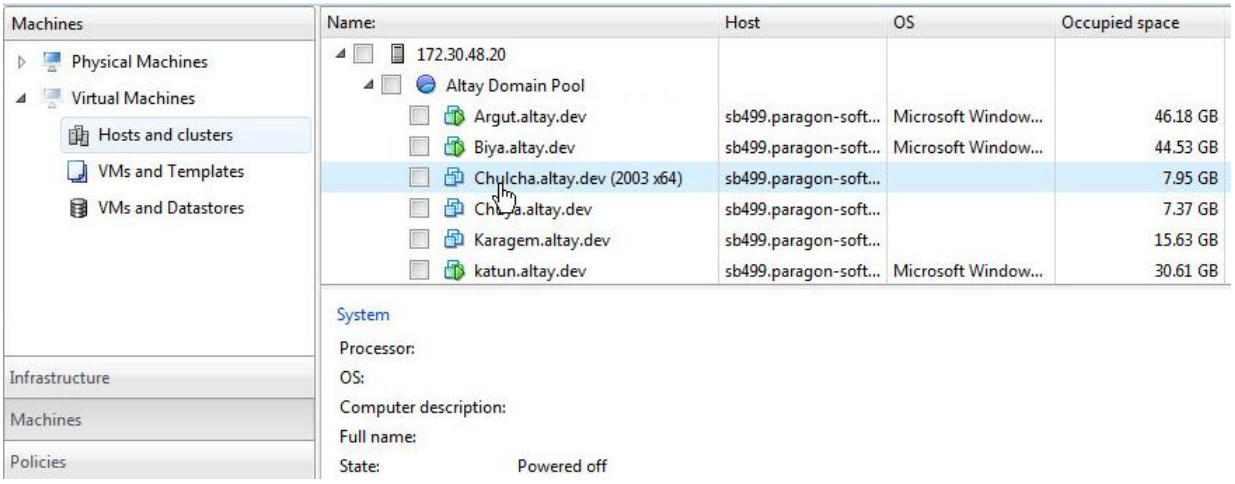

If ESX Agent has been installed, but no ESX host is displayed here, first click **Refresh** to force update of the information. If still no result, please check the following:

- A network adapter on the machine where ESX Agent is installed can access outer resources;

- The provided ESX access credentials are valid and allow enough privileges. You can always edit credentials by selecting **Infrastructure > ESX Agents > Edit**.

Editing ESX Connections

1. Select **Infrastructure > ESX Agents**, then click on the required ESX connection and select **Edit ESX connection**.

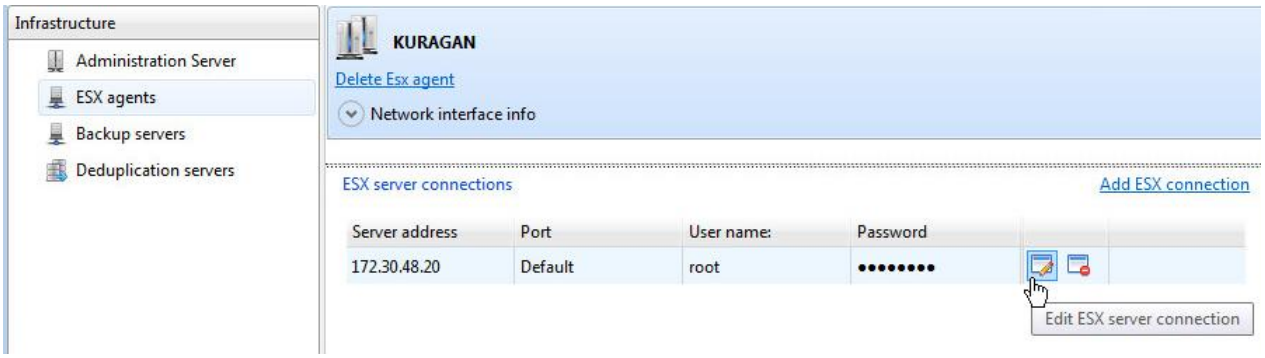

2. Modify all necessary parameters, then click **Save changes**.

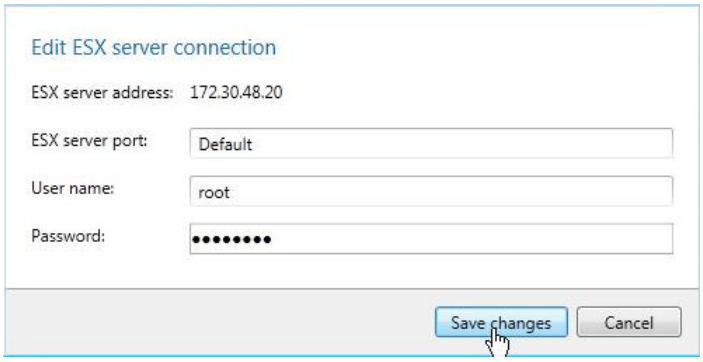

Deleting ESX connections

1. Select **Infrastructure > ESX Agents**, then click on the required ESX connection and select **Remove ESX connection**.

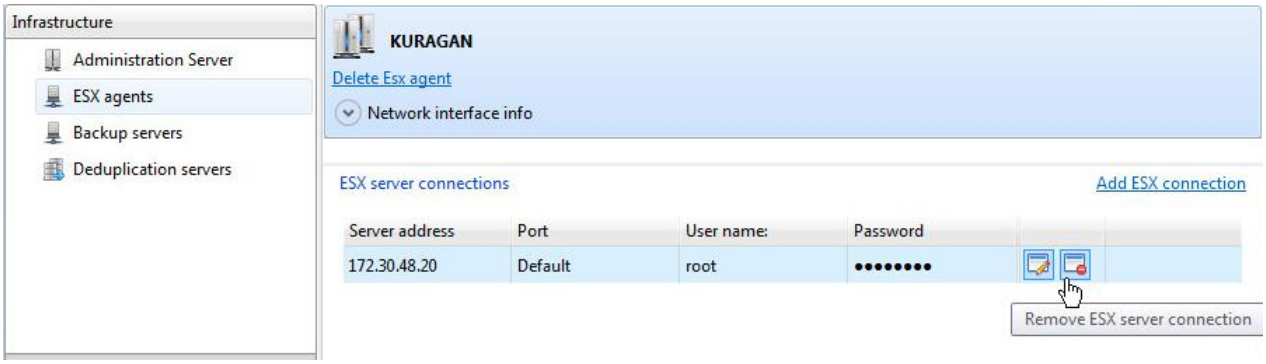

2. Confirm the operation.

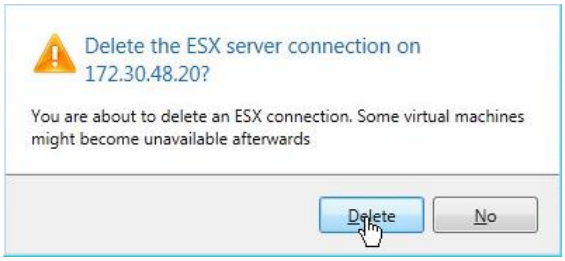

Managing Backup Server

Configuring Backup and Replica Storages

Backup Server supports two-tier backup storage architecture of primary and secondary storages, enabling to set up data migration policies (Storage Archiving) between them according to a particular disaster recovery plan. It's up to you to use this option or not, but if you want Backup Server to store backups of physical and/or virtual machines or VM replicas, you need to register at least one primary storage of the appropriate type.

The current version of the product can back up physical or virtual target machines either to a local folder of Backup Server or to a network share, while replicate virtual machines directly to an ESX datastore. Thus you should configure at least two primary storages if you'd like both to back up and replicate, which is easy to achieve since one Backup Server can manage several storages at a time.

It's allowed to configure several Backup Servers, each having several primary and secondary storages.

Registering primary storages

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Backup Server** is installed.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. If you've got several backup servers, first select the required backup server by clicking on its name, then click on the **New primary storage** link.

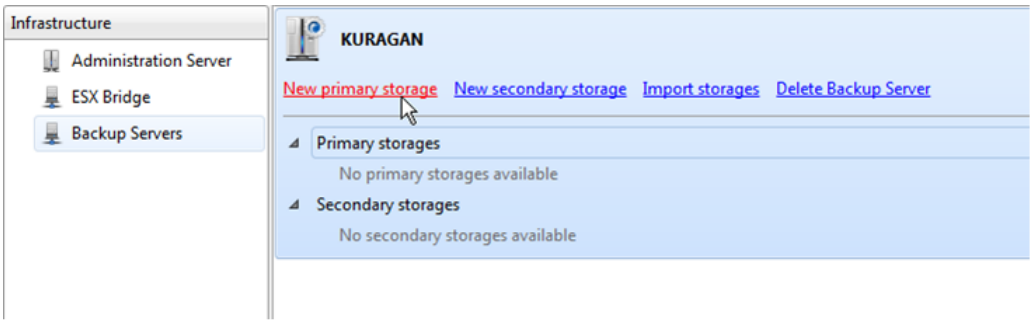

You can also initiate this operation by the right click of the mouse button, then selecting the corresponding option.

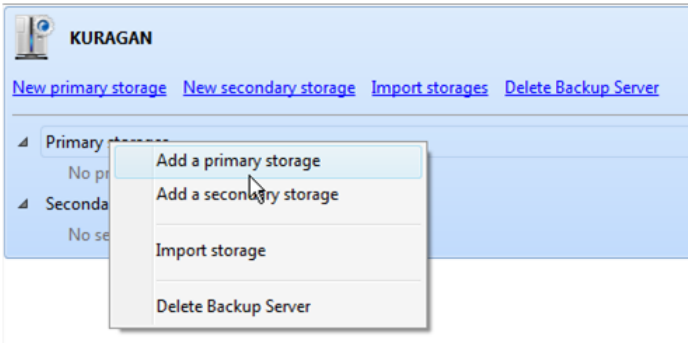

4. In the opened dialog you can see several tabs that enable to configure a primary storage of the required type:
   - **Common properties** to set the main parameters of a primary storage;
   - **Archiving** to create an archiving policy for it (please consult the Setting up a dual backup strategy chapter for more information);
   - **Free space notification** to make the program generate email notifications when free space on the storage is about to deplete;
   - **Deduplication** to link the storage to Deduplication Server (please consult the Linking Backup Storages to Deduplication Server chapter for more information);
   - **Security** to encrypt the storage against unauthorized access;
   - **Advanced settings** to modify advanced parameters (available if the corresponding option is active in the Settings dialog).

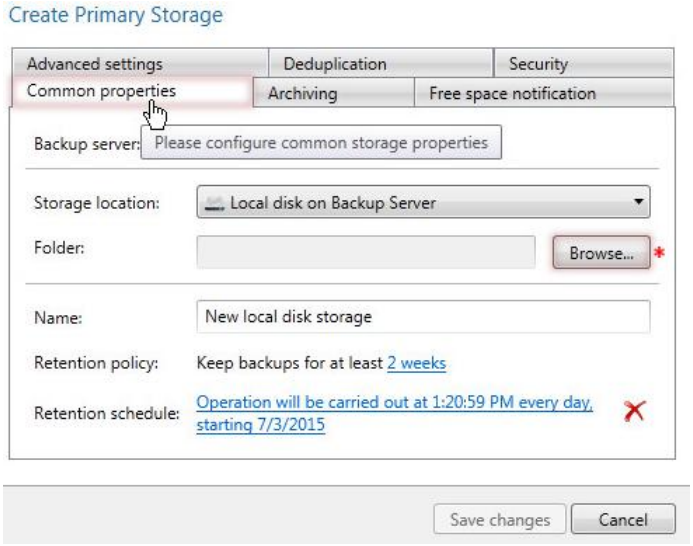

For local storages (can store backups of physical and virtual machines):

- **Storage Location**. Select **Local disk on Backup Server** for storing backups of target machines locally.

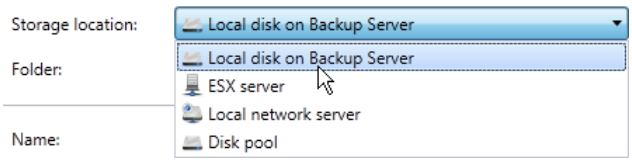

- **Folder**. Click **Browse** to specify a local disk and folder on Backup Server to place backup images to. Use the **New folder** button if necessary. Please make sure the amount of free space on the selected volume is enough to store all future backup images.

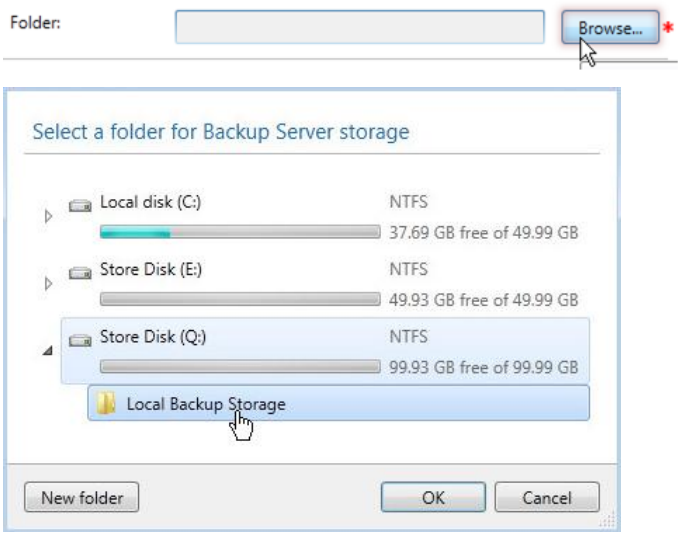

- **Name**. Give it a catchy name.

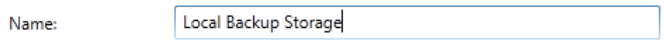

- **Retention policy**. Click on the default option (two weeks) to specify how long backups should be kept on the storage (from one second to infinite time) by using the appeared dialog. By default, for every machine our product creates a full backup during the first run, then incremental updates according to a set timetable. When time comes, all restore points beyond the set limit are merged with their full backup thus creating a new full backup. For more details please consult the Advanced settings section.

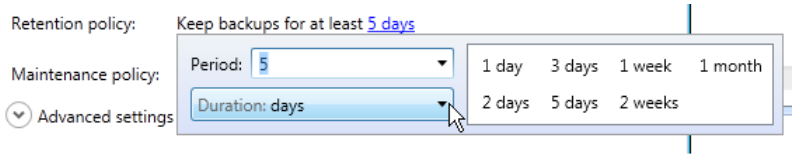

- **Retention schedule**. By default, all backup images inside the storage will be checked for and processed according to the specified data retention policy (general or individual) every day. If you want to change the default schedule, just click on the corresponding link.

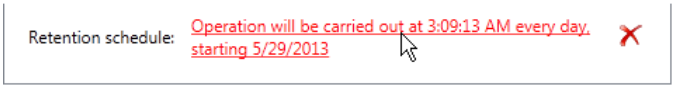

The opened dialog consists of two sections:

Basic scheduling

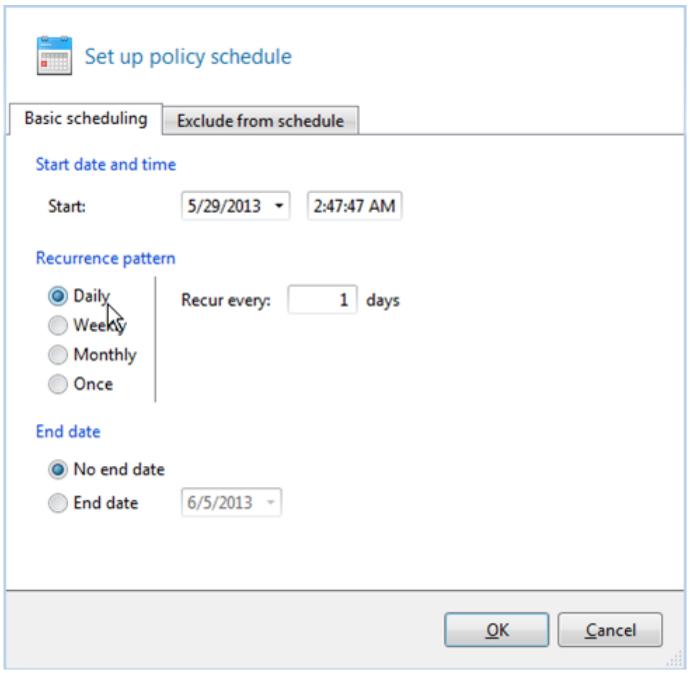

In this section you can set up a data retention timetable. The minimal available update interval is one day.

Exclude from schedule

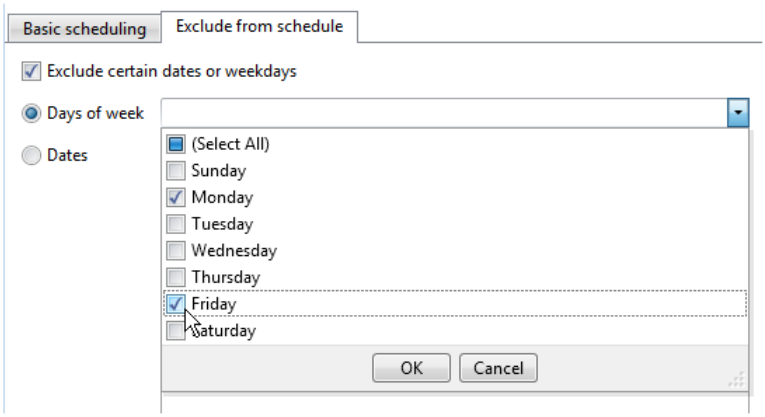

In this section you can specify days of week, or certain dates, when the retention policy should not be submitted.

Backup Server checks for and processes backup data of each storage according to their retention schedules. So if a retention policy limits the storage of backup images for two weeks for instance, while the specified retention schedule is configured to start up once a month, then backup images will only be checked for and processed once a month.

- Click on the **Free space notification** tab to set up email notifications when the storage reaches the specified free space limit (10 GBs by default). The corresponding warnings will also be generated in the Event Journal after each backup or maintenance policy.

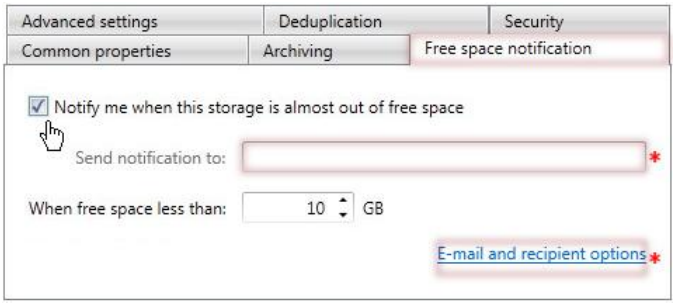

To learn how to configure the email transport system, please click here.

- Click on the **Security** tab to protect the storage against unauthorized access through an industry-standard 256-bit AES encryption algorithm.

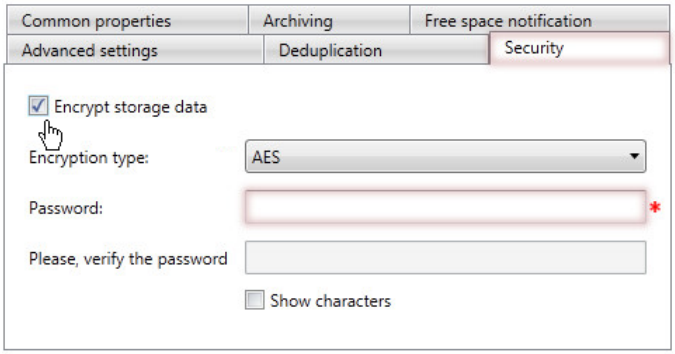

Encryption is available for disk pools, local and network (UNC) backup storages only.

- Click on the corresponding tab to see and configure additional options if necessary:

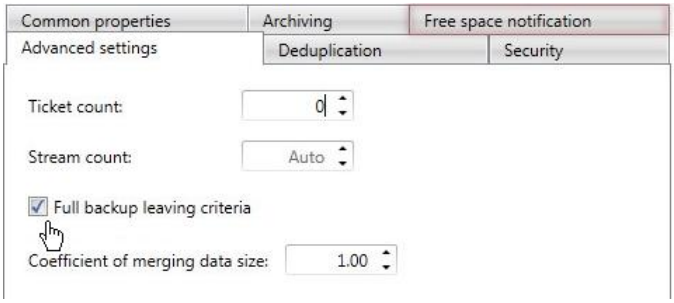

  - ***Ticket count***. Use this option to limit the number of simultaneously backed up machines to this storage (**Auto** by default). This can help when Backup Server lacks CPU, disk subsystem and/or network throughput performance;
  - ***Stream count***. Use this option to limit the number of simultaneously opened connections to backup objects (**Auto** by default);
  - ***Full backup leaving criteria***. This option can help to avoid time-consuming operations involving the merge of a big amount of data of a full backup with its relatively small increment, every time the storage data retention policy thins out backups. As a result only small increments are allowed to be merged. As you can see there's also the **Coefficient of merging data size** parameter set to 0.50 (50%), which is

actually specifies that a full backup will be allowed to merge if an increment coming after it is half of its size (30% - 50% are recommended values).

Thus use this option to suppress merge of a full backup to its increment (allowing increments to be merged only) when it becomes outdated according to the specified data retention policy, until a new full backup is created, or until an increment of the corresponding backup chain reaches the size specified in the merge data size criteria. Please note if you don't use this option when backing up a large amount of data (more than 100 gigabytes), please be ready to very poor performance of the data retention mechanism.

Advanced settings will be available if the corresponding option is enabled in the Settings dialog.

- Click **Save changes** to complete configuration of the backup storage.

For network storages (can store backups of physical and virtual machines):

- **Storage location**. Select **Local network server** for storing backups of target machines locally.

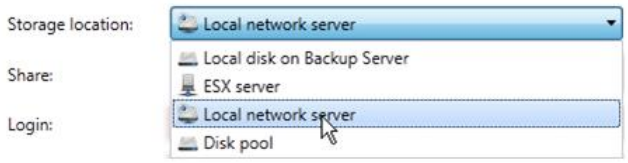

- **Share**. Specify the required network share by manually entering its location or click **Browse** to find it on the net.

Manually:

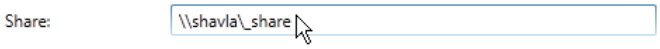

Through browsing:

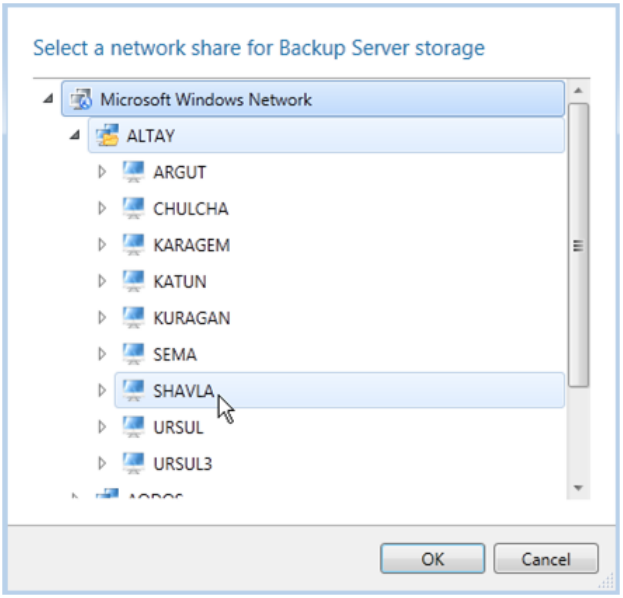

Double click on the required network machine to provide access credentials.

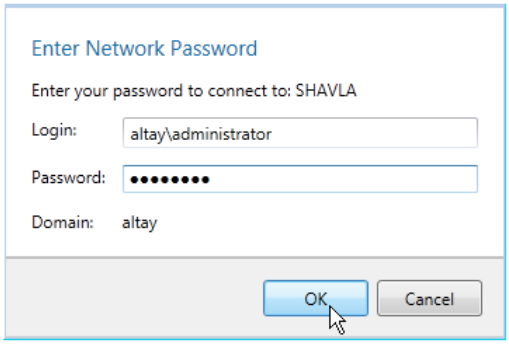

If the provided credentials are valid, you will be able to browse the specified network machine for the required storage folder. Click **OK** when ready.

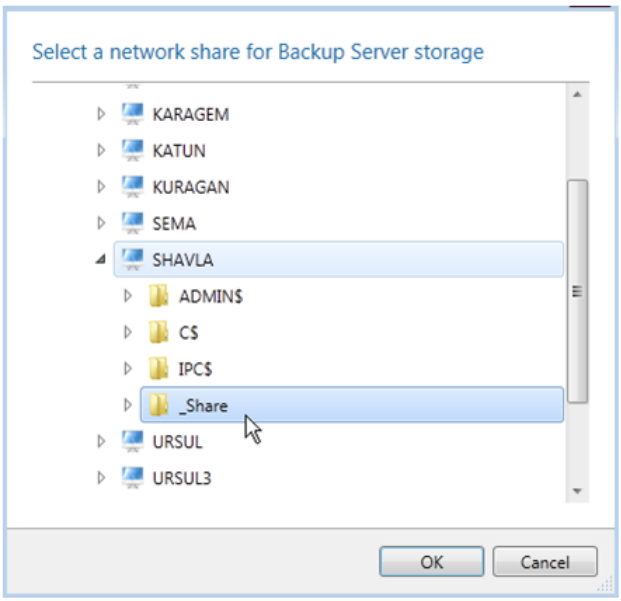

Backup to a network share located on a machine under control of a non-server OS may fail due to certain limitation on simultaneous connections. Thus if you're planning to store many backup objects (backup catalogs with multiple incremental updates), we strongly recommend you to use a specially -dedicated network repository for that.

- **Login and Password**. Specify access credentials for the manually provided network resource.

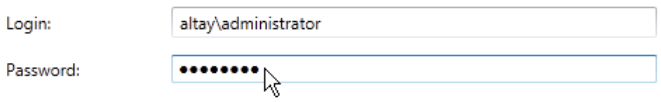

- **Name**. Give it a catchy name.

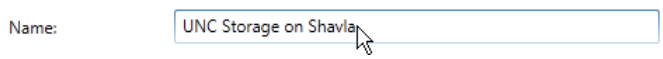

- **Retention policy**. Click here to know more on the subject.
- **Retention schedule**. Click here to know more on the subject.
- Encrypt the storage if necessary.
- Modify advanced settings if necessary.
- Click **Save changes** to complete configuration of the backup storage.

For ESX storages (can only store VM replicas):

- **Storage location**. Select **ESX server** for storing replicas of target machines.

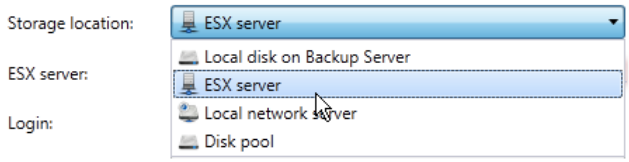

- **ESX server**. Click **Select storage location** to specify an ESX host, resource pool, and datastore to place replicas to. By default, there will be pre-selected parameters of the last added ESX connection. If the provided IP and access credentials are valid, there will be established connection to the specified VMware infrastructure.

If the required ESX host is a member of a vCenter, always use the IP address and credentials of that vCenter.

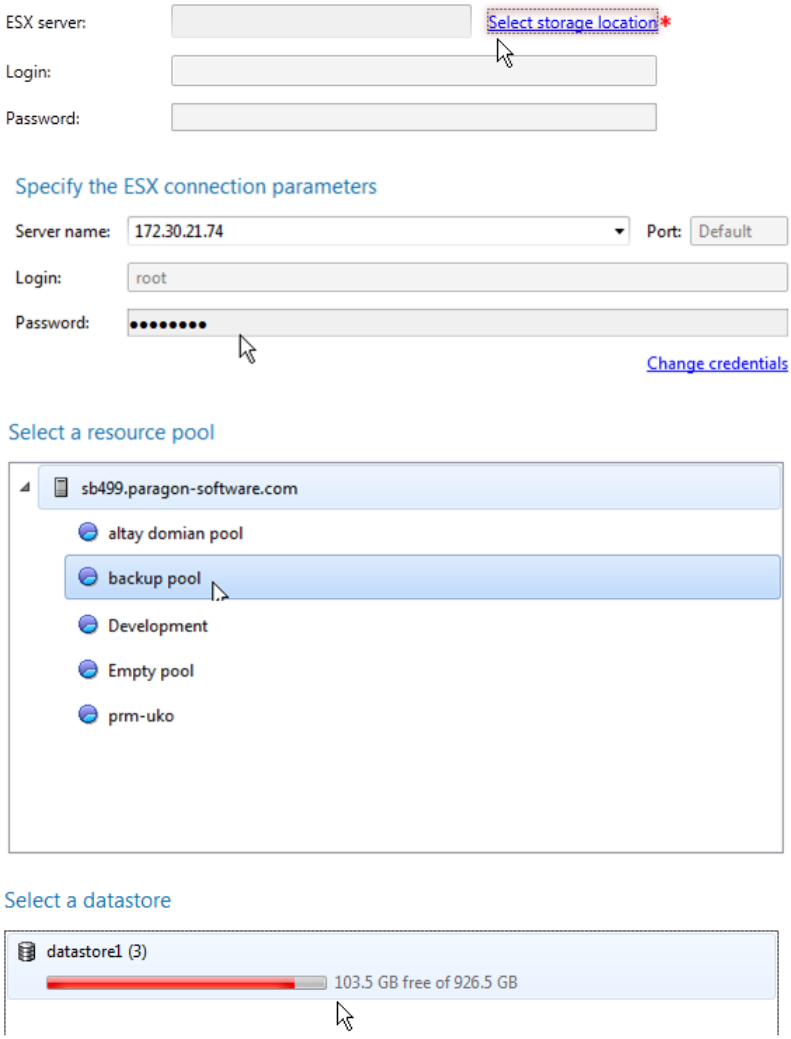

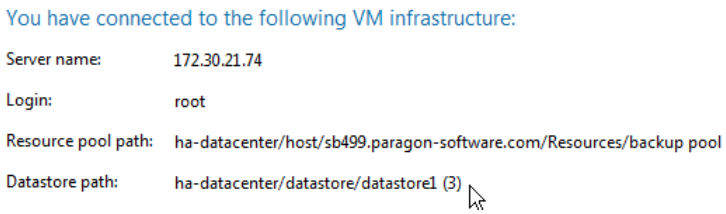

- **More ESX server connection info**. Click on the corresponding arrow button to see detailed information on the connected VMware infrastructure if necessary.

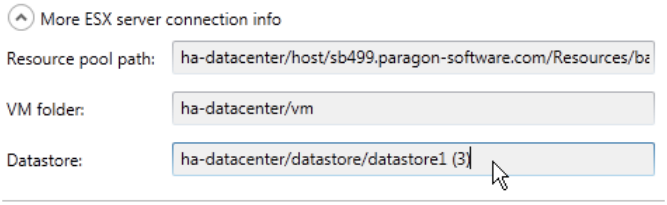

- **Name**. Give it a catchy name.

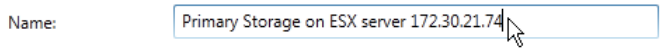

- **Retention policy**. Click here to know more on the subject.
- **Retention schedule**. Click here to know more on the subject.
- Modify advanced settings if necessary.
- Click **Save changes** to complete configuration of the backup storage.

For disk pool (can store backups of physical and virtual machines):

- **Storage Location**. Select **Disk Pool** for storing backups of target machines.

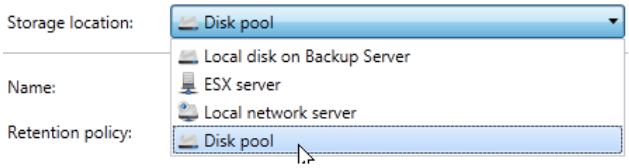

- **Name**. Give it a catchy name.

- **Retention policy**. Click here to know more on the subject.
- **Retention schedule**. Click here to know more on the subject.
- Click on the **Disk pool** tab, then the **Add new disk** link to add a disk to the pool. PPR enables to have an unlimited number of disks in one pool.

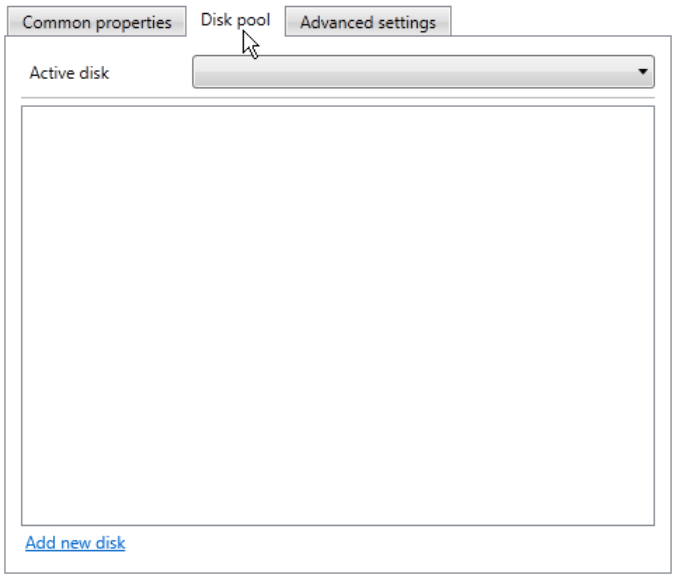

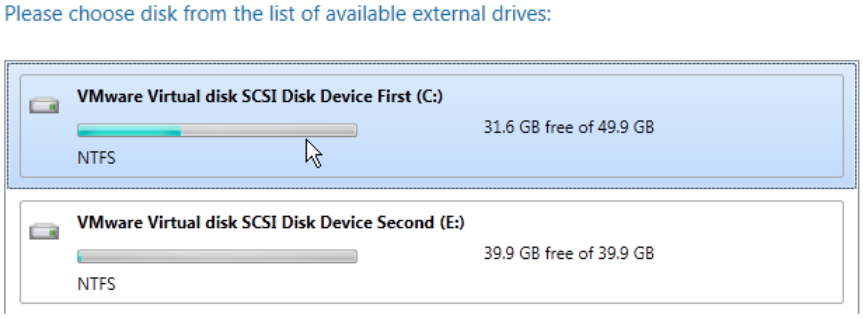

- Modify advanced settings if necessary.
- Click **Save changes** to complete configuration of the backup storage.

5. As a result you should have a new primary storage of the specified type and properties registered on Backup Server. If you need to register another primary storage, please go through this scenario once again.

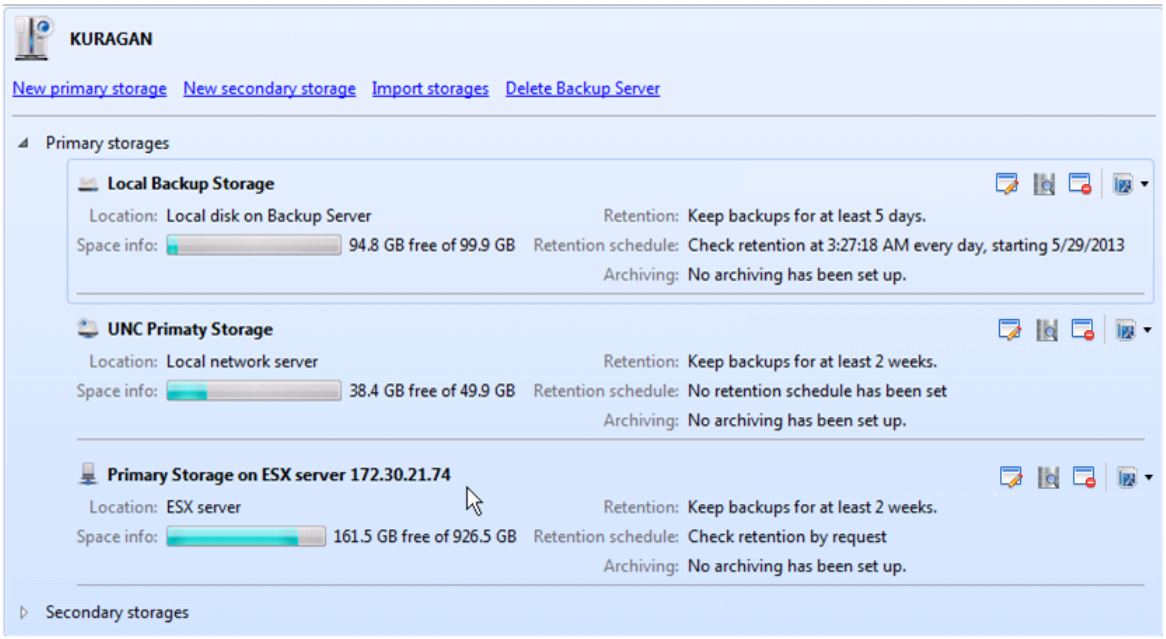

`

If you'd like to know how to manage storages and backup data they contain, please consult the corresponding chapters.

Registering secondary storages

Our approach to the two-tier backup storage infrastructure implies no direct use of second-tier (secondary) storages, but only through special archiving policies involving copying of data from primary storages by schedule or when backup is completed. Thus we don't need to do extra snapshots and load VMware ESX to accomplish parallel backup or replication of virtual machines. Moreover our approach allows easy on-the-fly conversion of backups to replicas or vice versa depending on the target secondary storage type.

In the current version of the product the following locations for residing a secondary storage are available:

- A local folder of a backup server where you're attempting to register the storage (can only contain backup images);
- An ESX datastore (can only contain replicas);
- A network share (can only contain backup images).
- An FTP online storage (can only contain backup images).

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Backup Server** is installed.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. If you've got several backup servers, first select the required backup server by clicking on its name, then click on the **New secondary storage** link.

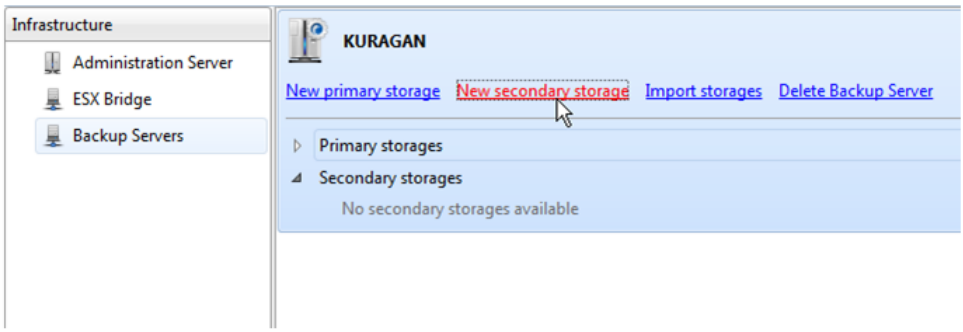

You can also initiate this operation by the right click of the mouse button, then selecting the corresponding option.

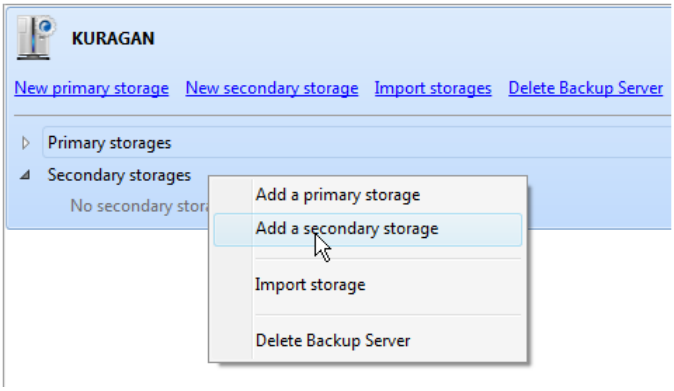

4. In the opened dialog configure a secondary storage of the required type:

   For local storages (can store backups of physical and virtual machines):

> **The process of configuring a secondary local storage is identical to that of the primary one. Click here to know more on the subject.**

   For ESX storages (can only store replicas of virtual machines)**:**

> **The process of configuring a secondary local storage is identical to that of the primary one. Click here to know more on the subject.**

   For network storages (can store backups of physical and virtual machines):

> **The process of configuring a secondary local storage is identical to that of the primary one. Click here to know more on the subject.**

   For FTP storages (can store backups of physical and virtual machines):

   - **Storage location**. Select **FTP site** for archiving backup or replica data to an online FTP storage.

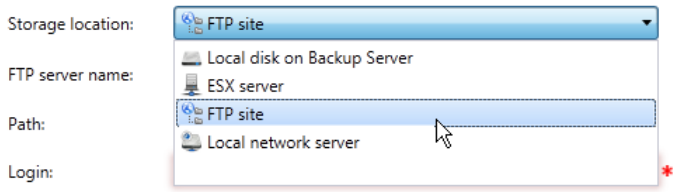

   - **FTP server name and Port**. Type in an address of the desired server and specify the required communication port (21 for FTP by default).

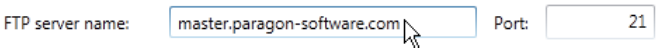

> **You need to check out yourself Windows Firewall or programs of this kind allow PPR to work with the required port.**

   - **Path**. Type in an exact location for archived data.

   - **Login and Password**. Provide access credentials.

   - **Name**. Give a catchy name to the FTP connection.

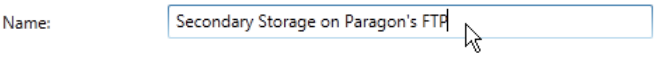

No data retention is supported for FTP storages, thus you should manually manage backup data they contain through the Browse Storage function.

- Modify advanced settings if necessary.
- Click **Save changes** to complete configuration of the backup storage.

5. As a result you should have a new secondary storage of the specified type and properties registered on the selected backup server. If you need to register another secondary storage, please go through this scenario once again.

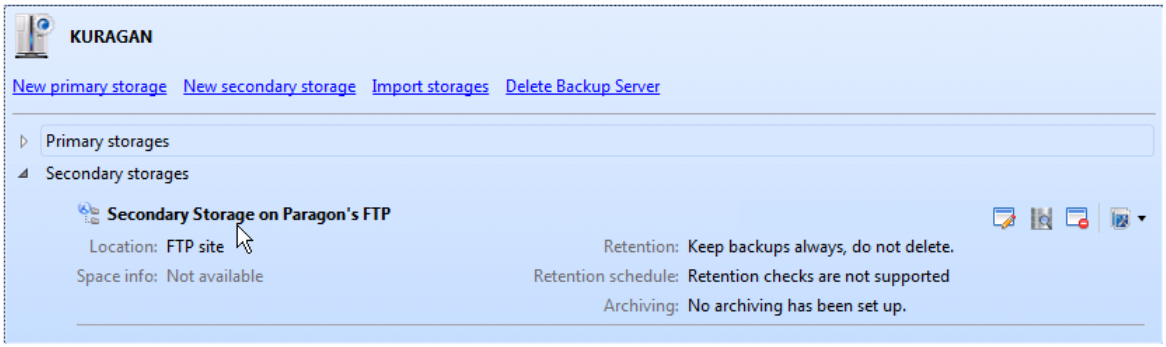

Setting up a dual backup strategy

Our approach to the two-tier backup storage infrastructure implies no direct use of secondary storages, but only through special archiving policies involving transfer of backup data from primary to secondary storages by schedule. Unlike parallel backup techniques, we enable to:

- Avoid extra snapshots of ESX guests or physical Windows machines. Thus we minimize backup windows, releasing CPU and network bandwidth resources.
- Do on-the-fly conversion of backups to replicas or vice versa depending on the target secondary storage type.
- Schedule transfer of backup data to secondary storages when most appropriate.

It's also allowed to configure archiving policies for secondary storages.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be registered at least one primary backup or replica storage.
- There should be registered a secondary backup or replica storage.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. On the right pane select the required storage (primary or secondary), then click the **Modify storage settings** icon or double click the storage.

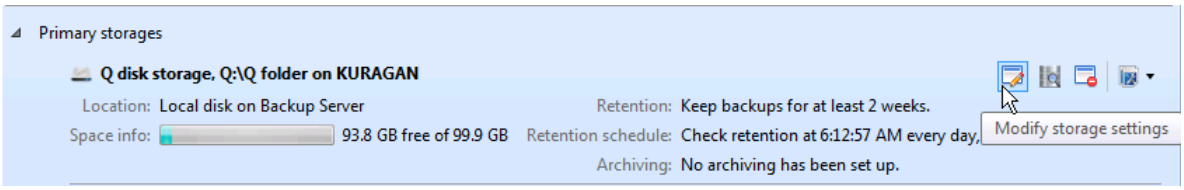

4. In the opened dialog click on the **Archiving** tab to specify the following parameters:

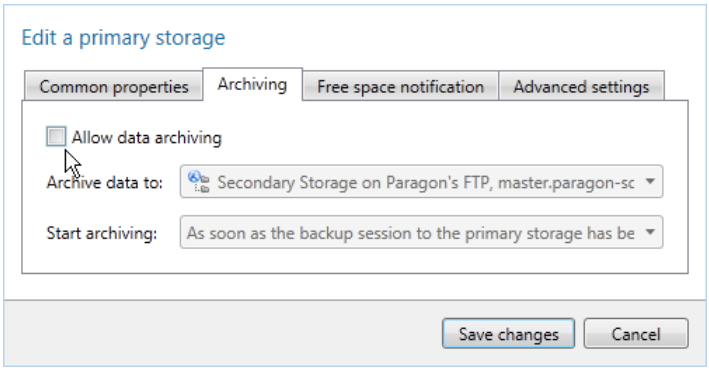

- **Allow data archiving**. Mark the checkbox to enable the data archiving mode.
- **Archive data to**. Select a target secondary storage (if several), where backups or replicas will be archived to.

If the chosen secondary storage is local, network or FTP, the resulted data it contains will be backup images, no matter which type of the primary storage is used, thus replicas from a primary ESX storage will be converted to backup images during the process.

If the chosen secondary storage is ESX, the resulted data it contains will be replicas, no matter which type of the primary storage is used, thus backup images from a primary local storage will be converted to replicas during the process.

- **Start archiving**. By default, archiving will be triggered once a backup session on the primary storage has been completed, but you can schedule the operation as well.

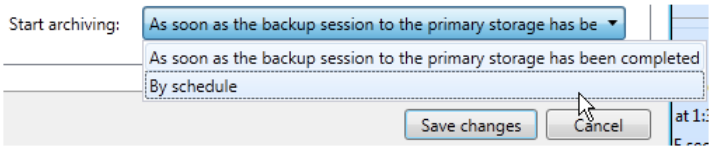

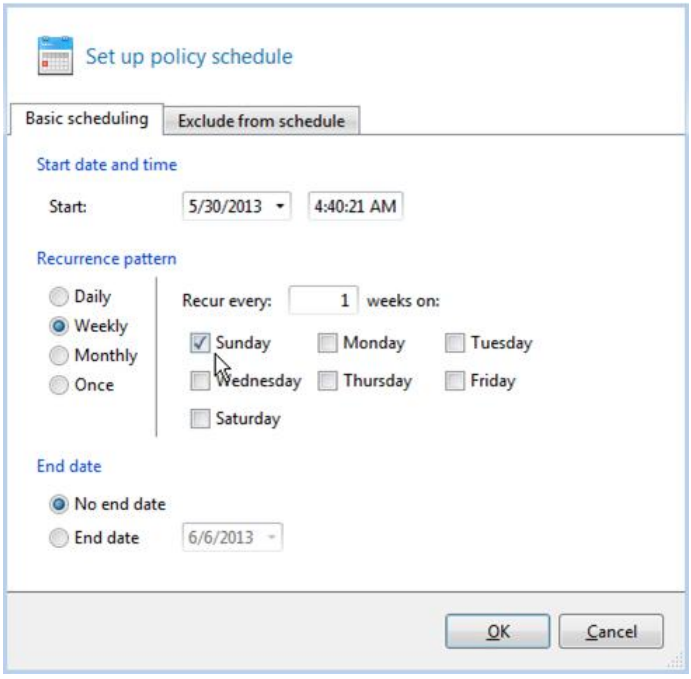

5. When ready, click **Save changes**. There will be created, and then validated a corresponding archiving policy, which you can see through a popup window.
6. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
7. You can either disable/enable or force launch of any archiving policy by using the corresponding icons.

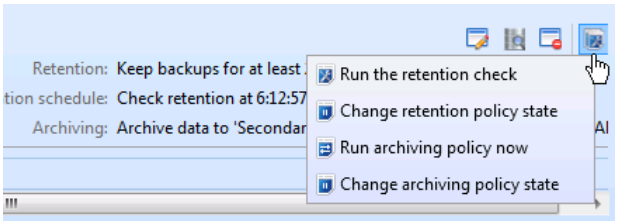

Data archiving is only applied to backup sessions that have not been archived yet. Thus if some backup sessions have been deleted on secondary storage manually of due to activity of a storage retention policy, these sessions won't be archived again.

Manual deletion of backup sessions on primary storage or their automatic deletion due to activity of a storage retention policy doesn't apply to secondary storage.

Treating backup data inconsistency on secondary storage

If there has been detected a problem with backup data consistency on secondary storage, a failed backup catalog(s) can be marked for repeated archiving from primary storage. It works this way – marked backup catalogs on secondary storage are automatically renamed by adding to their names the suffix **_MarkedInvalid**, and then during the next startup of the archiving policy corresponding backup data are archived from primary storage once again. Thus backup data from marked catalogs can be manually accessed for restore, deletion, etc.

Backup catalogs (a backup catalog contains all backup images of a particular target machine) that have problems with data archiving are highlighted in Storage Browser. Just click on a problem catalog to see its current archiving state.

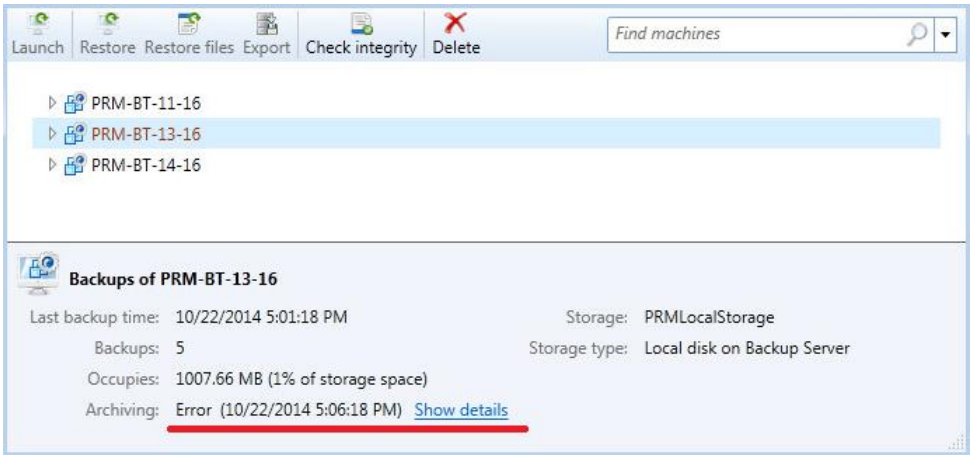

If you'd like to get more details, then click on the corresponding link. In the opened dialog you can see where the selected backup catalog was archived to and its archiving state. Besides, you can mark the catalog here for re-archiving.

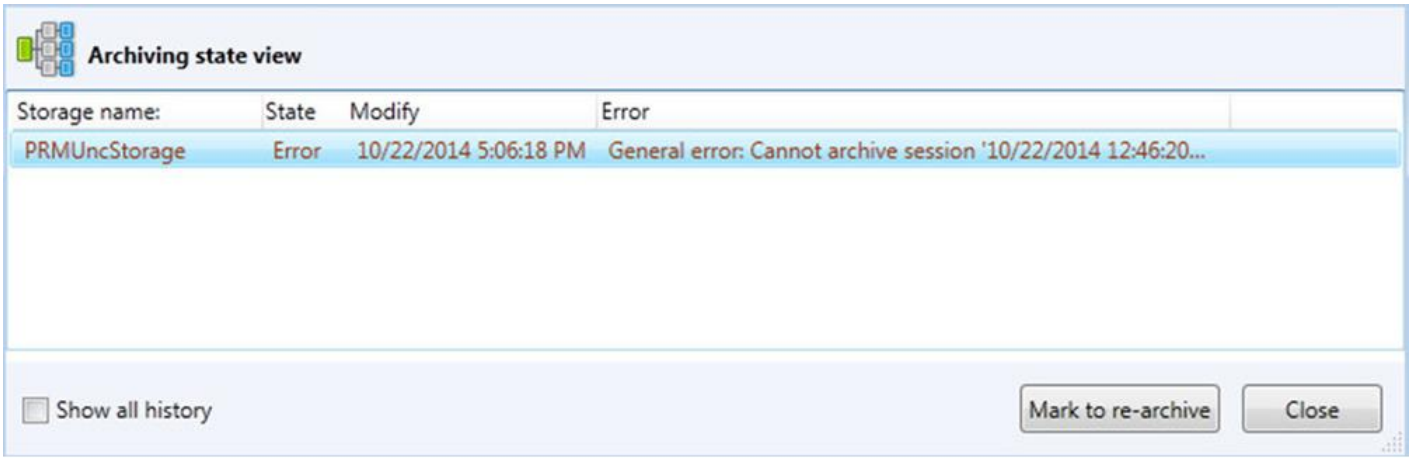

Each backup catalog is processed independently from others. Thus when encountering problems with archiving backup sessions of one machine, it doesn't influence the archiving of backup sessions that belong to other machines.

Prior every launch of an archiving policy, the latest recovery point on secondary storage for each machine is verified through mounting. This is the easiest way to make sure backup images are not corrupted. If the mount test ends with an error, the archiving policy ends with an error as well, while the problem recovery point is marked as invalid.

If the integrity checkup finds corrupted backup sessions on secondary storage, it marks them as invalid (Invalid flag). All backup sessions on secondary storage are checked for the presence of the Invalid flag during data archiving, and if found, there are two options:

- **If an invalid backup session belongs to the latest incremental chain (i.e. after creation of the latest full backup image), then the archiving policy ends with an error.**
- **If an invalid backup session belongs not to the latest incremental chain (i.e. before creation of the latest full backup image), then the archiving policy works normally and the latest recovery point will be ok. However, the archiving policy ends with a warning to notify the user that there are problems with backup data.**

Storage Maintenance

Modifying storage properties

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be configured at least one backup or replica storage.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, select **Infrastructure > Backup Servers**.
3. On the right pane select the required storage, then click the **Modify storage settings** icon or double click the storage.

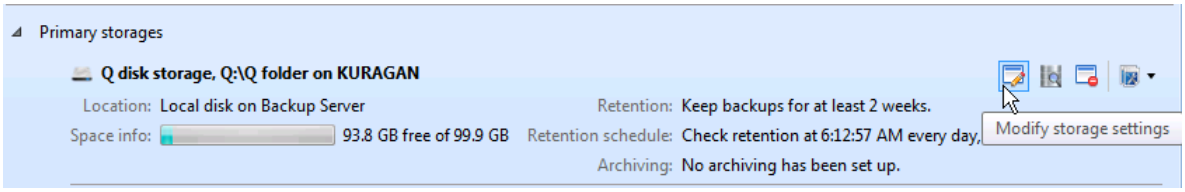

4. Change parameters according to your needs. Click here to know more on the subject.

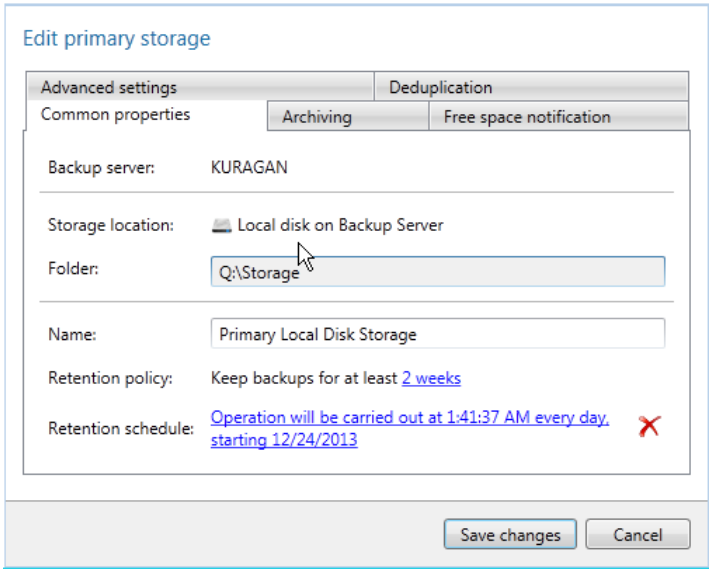

5. Click **Save changes** to complete configuration of the backup storage.

Administering storage backup data

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be configured at least one backup or replica storage containing a backup or replica.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, select **Infrastructure > Backup Servers**.
3. On the right pane select the required storage, then click the **Browse storage** icon.

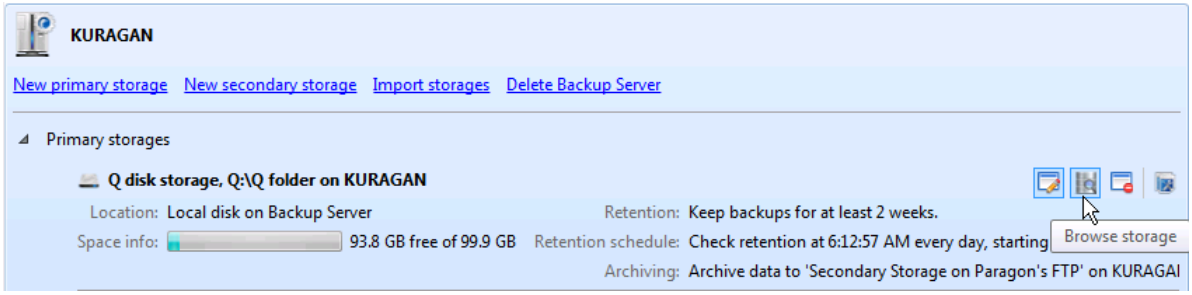

4. In the opened dialog you can see a list of all machines ever protected by PPR with available backups or replicas (depends on the storage type) of the specified storage. Use the **Find machines** field to quickly find the required object.

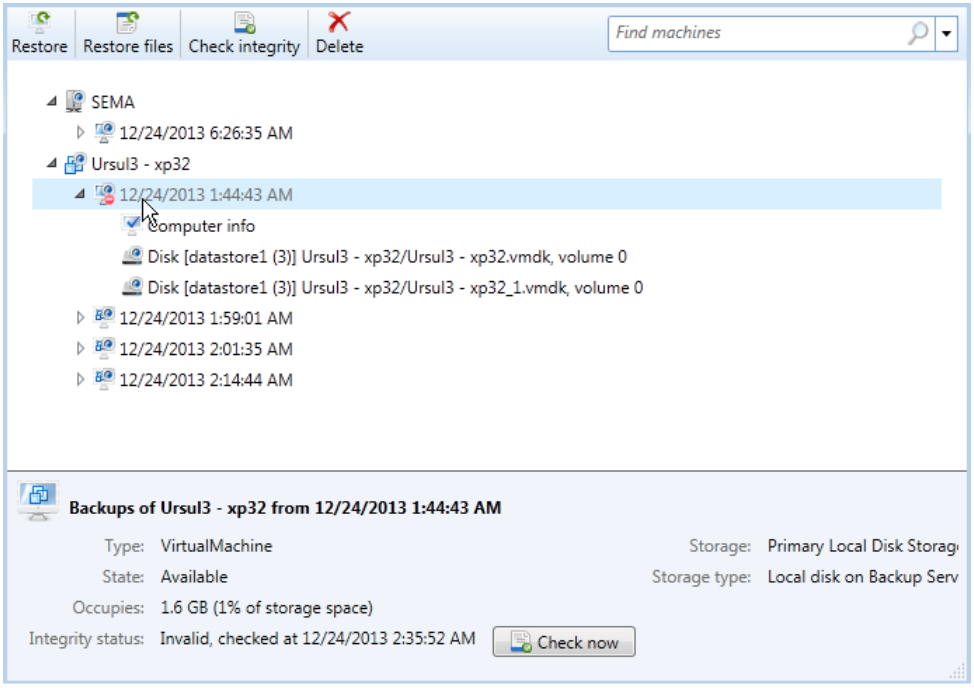

5. Select the required machine, then one of available restore points. Below you can see detailed information on the selected object. Use icons above to:
   - Initiate restore of the selected backup or launch of the selected replica. To know more on the subject, please consult restore scenarios.
   - Initiate restore of particular files and/or folders from the selected backup image or replica. To know more on the subject, please consult the Restoring Separate Files scenario.
   - Verify the selected object for errors by submitting a corresponding policy. If it turns out to be invalid, it will be marked by a special icon. We highly recommend you to check backup images for integrity before initiating restore.
   - Delete the selected object from the storage. Please note if you delete an increment from somewhere in the middle of the incremental chain, the program will automatically initiate a data merging operation to keep the incremental chain valid. This action may take some time (depends on the amount of data of the deleted object) during which all members of the corresponding incremental chain will stay unavailable (grey out).

Despite the fact that you're allowed to initiate complete restore or retrieval of certain files/folders from invalid backup images, please do it at your own risk.

Managing retention policies

You can either disable/enable or force launch of any retention policy.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be configured at least one backup or replica storage.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. On the right pane select the required storage, then click on the corresponding icon to do the required action (force launch of the retention policy in our case).

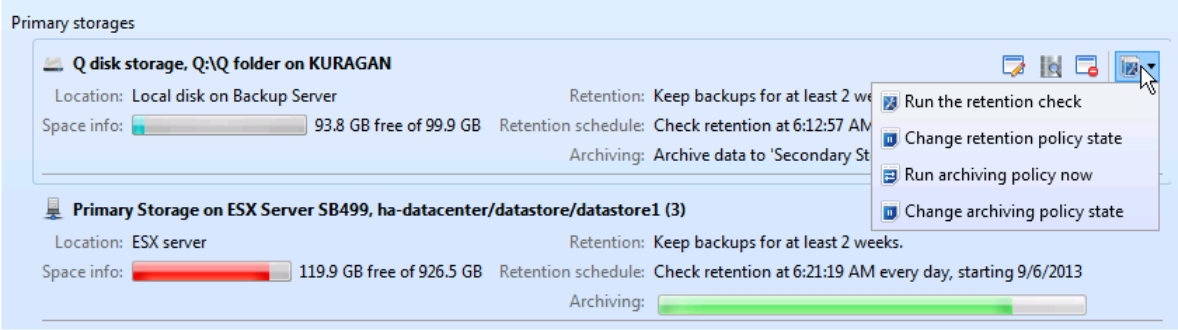

4. The retention policy will be submitted immediately, which you can see through a popup window.
5. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.

Attaching storages

You can attach to the infrastructure storages from another PPR infrastructure or those you've deleted previously to use backup data they contain. You've got several options:

- Configure a new backup server or use an existing one on a machine that contains the required storage, then attach this storage as a local backup storage.
- Use one of the registered backup servers to attach the required storage as a network backup storage.
- Configure a new backup server or use an existing one to attach ESX storage containing replicas of virtual machines.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- You should have a local, network, or ESX unregistered storage or mini-storage (Paragon's container aka pVHD type only) obtained during the Export backup operation that you're going to attach to the infrastructure. An unregistered storage is storage from another PPR infrastructure or that that you deleted from the infrastructure.

By default, data retention, archiving and other auxiliary options are reset to prevent untimely deletion of backup data and other problems. You can enable these functions through the 'Modify storage settings' option. Click here to know more on the subject.

PPR automatically detects the type of storage (primary or secondary) it's dealing with

allowing the PPR Administrator to choose the preferred type during attachment.

Attachment of registered storages is not recommended. To avoid any conflicts first delete a registered storage from another PPR infrastructure, and then repeat the operation.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. If you've got several backup servers, first select the required backup server by clicking on its name, then click on the **Attach storage** link.

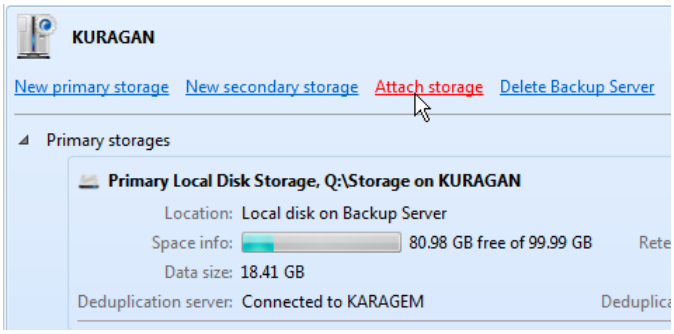

You can also initiate this operation by the right click of the mouse button, then selecting the corresponding option.

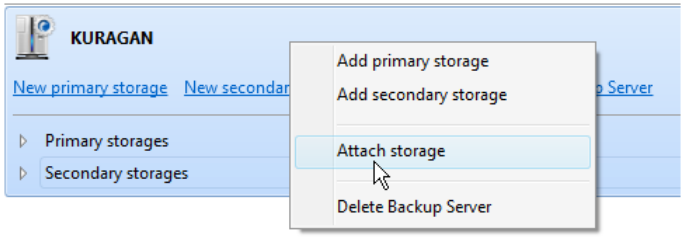

4. First you need to specify type of the storage you're going to attach (local, network, or ESX).

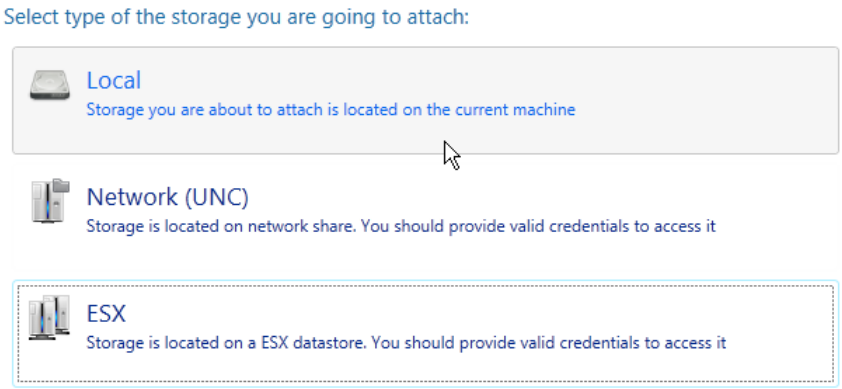

- **Local**. If the storage to attach resides on Backup Server you should set:
  - **Folder**. Click **Browse** to specify a local disk and folder on Backup Server where the required storage is located. PPR uses guid names for backup catalogs and machines, but backup storages are named traditionally, thus you can easily find what you need. If you know you've got several storages on a volume, mark the **Advanced search** option to make the wizard look for storages in all subfolders, otherwise only the specified location will be processed.

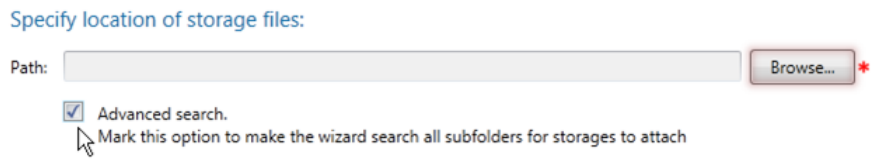

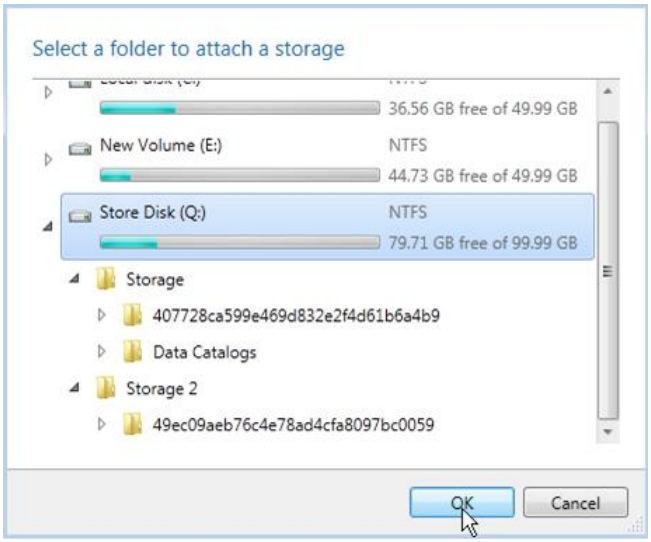

If there are several storages in the specified location, the wizard informs you about it, prompting to choose one of them to proceed. To help you make the right choice it also outputs a number of storage properties at this stage.

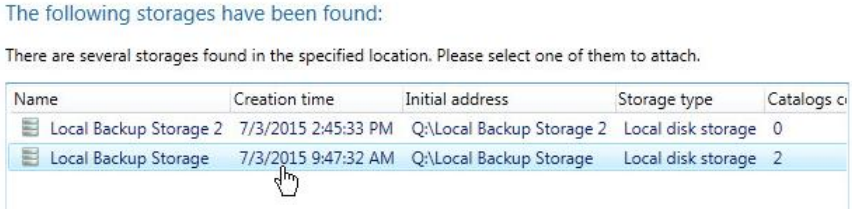

You will need to provide a password when attempting to attach an encrypted backup storage from another PPR infrastructure.

If the specified storage contains deduplicated backup sessions, you will be prompted to select an appropriate Deduplication Server.

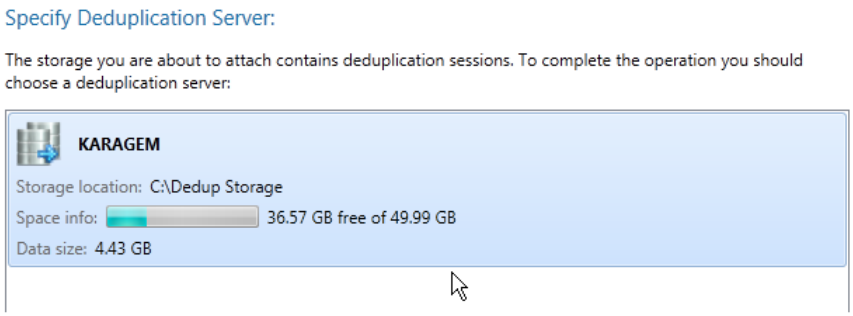

- **Additional options**. By default, the wizard offers to use the old storage name, which you can change to any of your choice. Besides here you can set the desired level of the integrity checkup for backup data the storage contains, change the original storage type, and reconfigure the data retention policy.

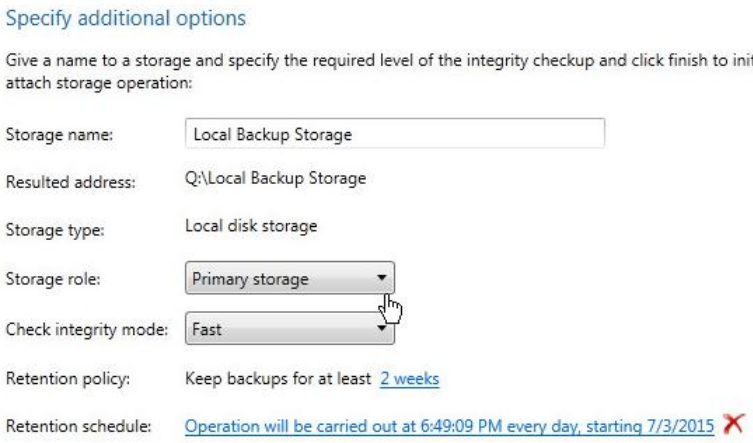

- **Network**. If the storage to attach resides on a network share you should set:
  - **Path**. Specify the required network share and access credentials manually in the corresponding fields or click **Browse** to find it on the net.

    By clicking on a network machine, the wizard attempts to access it. At this stage you should provide access credentials. If you know you've got several storages on a volume, mark the **Advanced search** option to make the wizard look for storages in all subfolders, otherwise only the specified location will be processed.

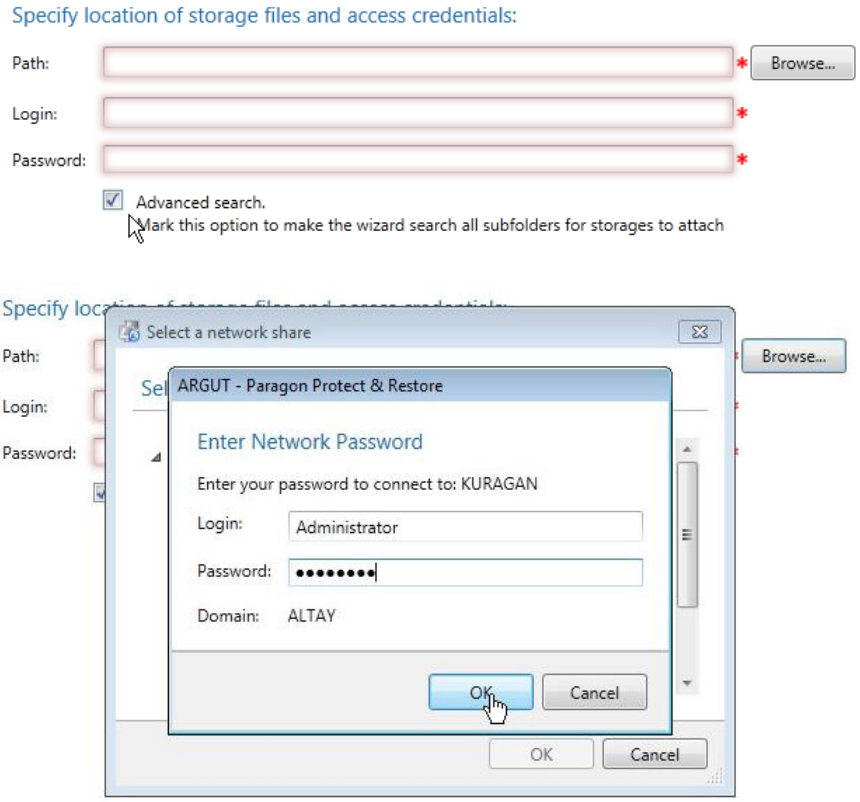

If the provided credentials are valid, you will be able to browse the specified network machine for a storage folder. Click **OK** when ready.

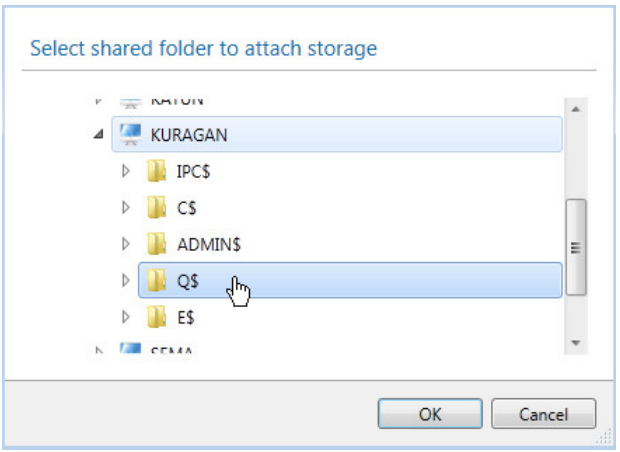

If there are several storages in the specified location, the wizard informs you about it, prompting to choose one of them to proceed. To help you make the right choice it also outputs a number of storage properties at this stage.

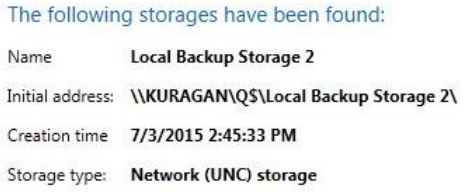

You will need to provide a password when attempting to attach an encrypted backup storage from another PPR infrastructure.

If the specified storage contains deduplicated backup sessions, you will be prompted to select an appropriate Deduplication Server.

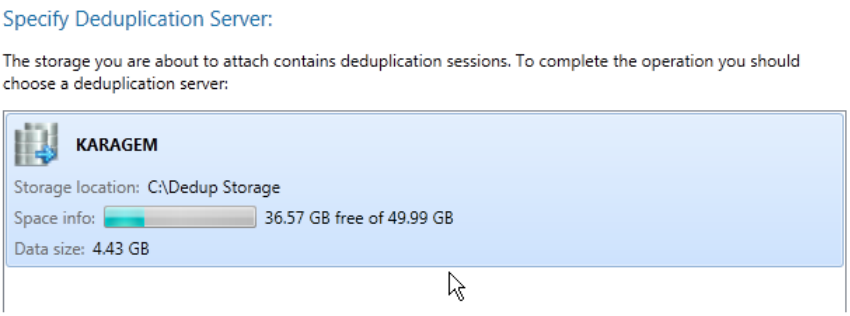

- **Additional options**. By default, the wizard offers to use the old storage name, which you can change to any of your choice. By default, the wizard offers to use the old storage name, which you can change to any of your choice. Besides here you can set the desired level of the integrity checkup for backup data the storage contains, change the original storage type, and reconfigure the data retention policy.

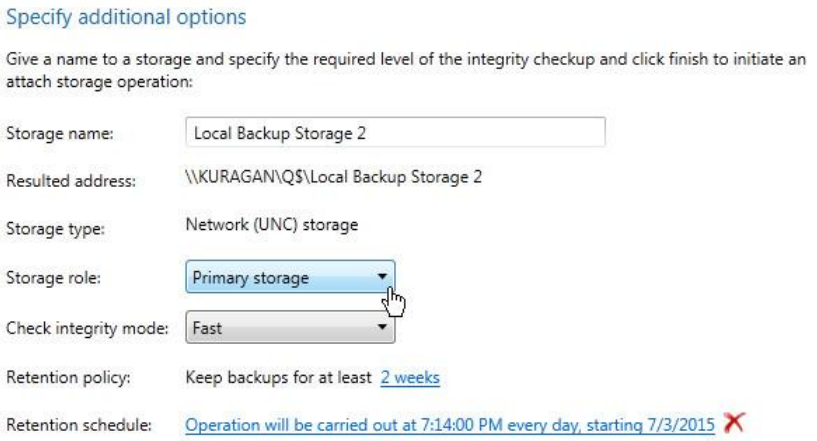

- **ESX**. If the storage to attach resides on a ESX datastore you should set:
  - **ESX connection parameters**. Enter a DNS name or IP address of the required vCenter or ESX host, a communication port (if necessary), and administrator credentials in the corresponding fields. Click **Save changes** to start the operation, which takes approx. a couple of seconds. If a success, the wizard will browse the specified resource.

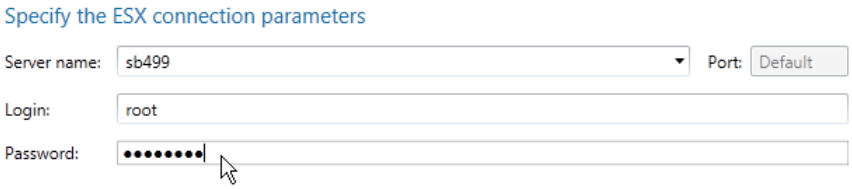

If the required ESX host is a member of a vCenter, always use the IP address and credentials of that vCenter.

  - **Resource pool**. Specify a resource pool that contains storage you'd like to attach.

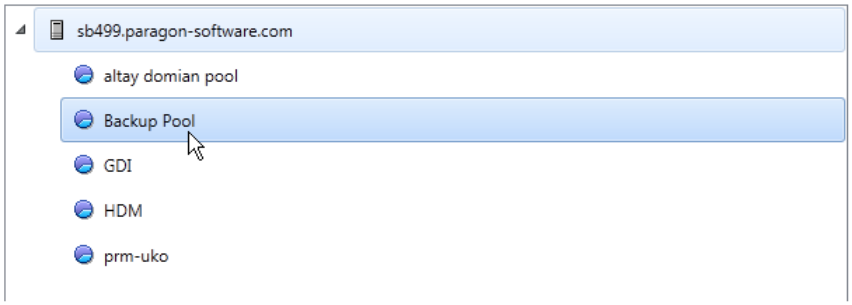

  - **Datastore**. Select a datastore where the required storage resides.

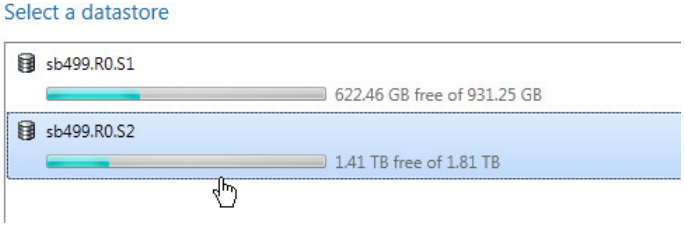

If there are several storages in the specified location, the wizard informs you about it, prompting to choose one of them to proceed. To help you make the right choice it also outputs a number of storage properties at this stage.

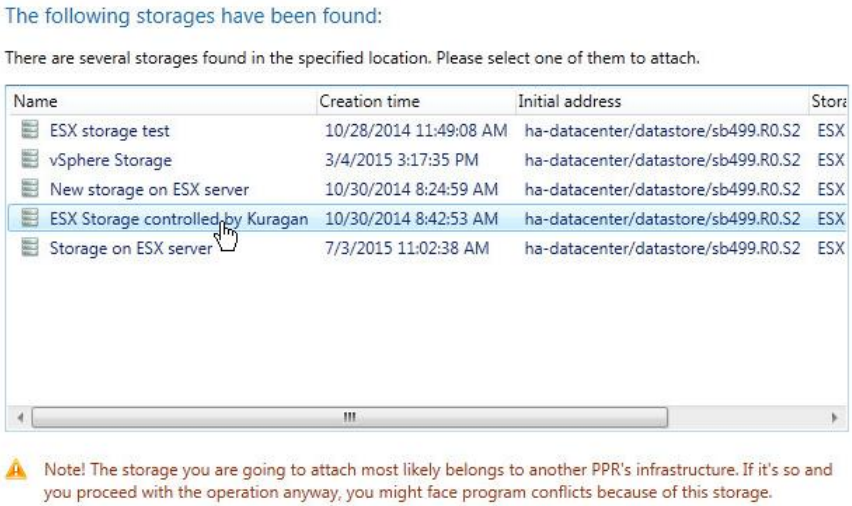

Attachment of registered storages is not recommended. To avoid any conflicts first delete a registered storage from another PPR infrastructure, and then repeat the operation.

- **Additional options**. By default, the wizard offers to use the old storage name, which you can change to any of your choice. Besides here you can set the desired level of the integrity checkup for backup data the storage contains, change the original storage type, and reconfigure the data retention policy.

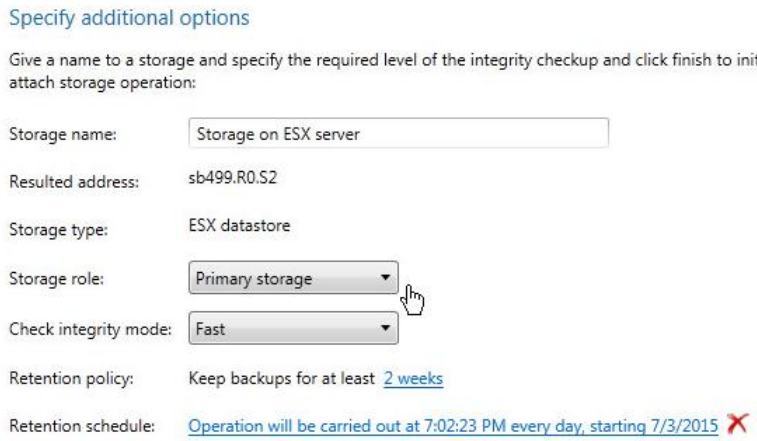

5. Click **Finish** to complete configuration of the backup storage. The operation will be initiated immediately, which you can see through a popup window.
6. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
7. When the attach storage operation that involves two actions (storage rebuild and data integrity verification) is over, its status will be updated. As a result you should have a new primary or secondary storage of the specified type (local, network, or ESX) and properties registered on Backup Server.

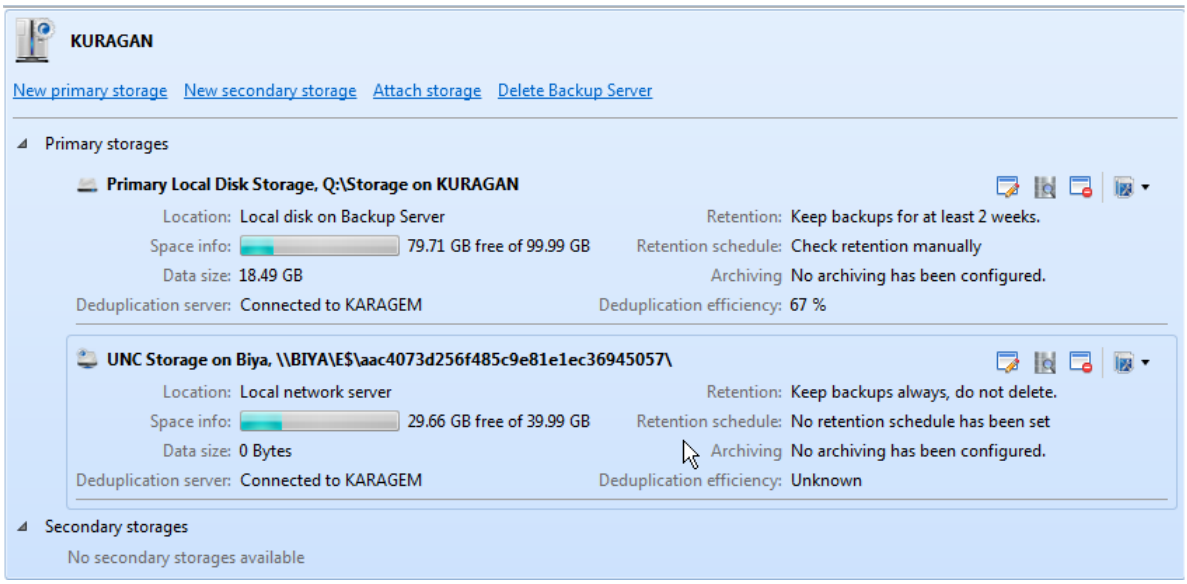

`

If you'd like to know how to manage storages and backup data they contain, please consult the corresponding chapters.

Deleting storages

You can unregister (delete) any backup or replica storage from the infrastructure. This operation doesn't involve deletion of backup data of the specified storage, but only its removal from the infrastructure. Later you can register this storage again through the Attach Storage function. If you'd like to delete backup data, please use the Browse Storage function first.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be configured at least one backup or replica storage.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. On the right pane select the required storage, then click the **Delete storage** icon.

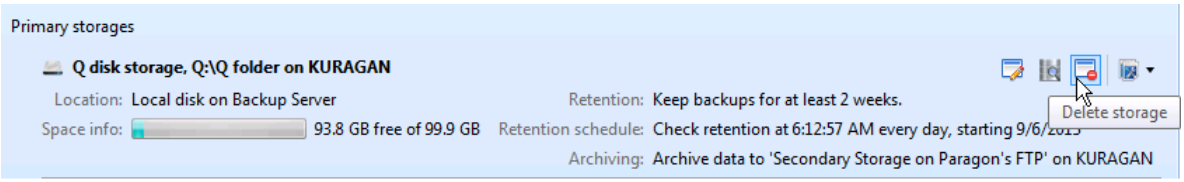

4. In the opened dialog check what protection policies are using this storage by clicking on the arrow button. If you're ok that all listed here policies will be disabled and their backups will become not available to use, then click **Delete** to confirm the operation.

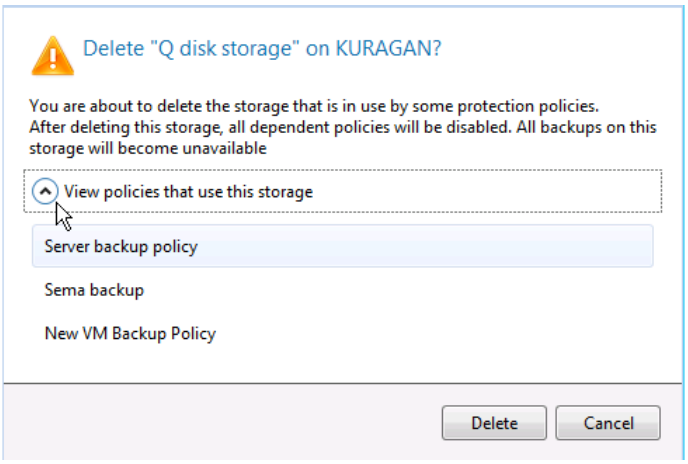

5. The operation will be initiated immediately, which you can see through a popup window.
6. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.

Exporting individual restore points

You've got the option to export any restore point from the PPR infrastructure to do on-site recovery of a certain machine through the WinPE media even when Administration Server or Backup Server(s), or both are damaged or not available on the net. During the export operation, there will be created a mini-storage containing the specified restore point in the required format (Paragon's container aka pVHD, VMWare VMDK, VHD, VHDX, or VirtualBox VDI). If having to do with an incremental or/and deduplicated image, PPR will detect it and automatically reconstruct it to a non-deduplicated full image.

In case of emergency, you can start up the target machine from the WinPE media, attach the resulted mini-storage (pVHD type only), and accomplish restore.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be configured at least one backup or replica storage containing a backup or replica.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Export backups**,

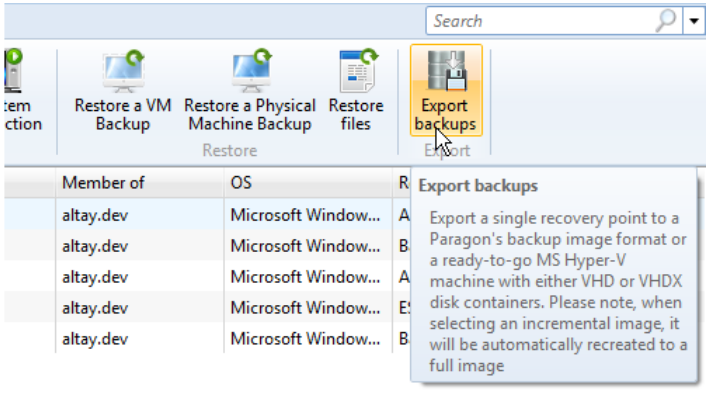

or go to **Infrastructure > Backup Servers**, select the required storage, then click the **Browse storage** icon to see machines it contains. Select the required restore point, then click **Export** to initiate the operation.

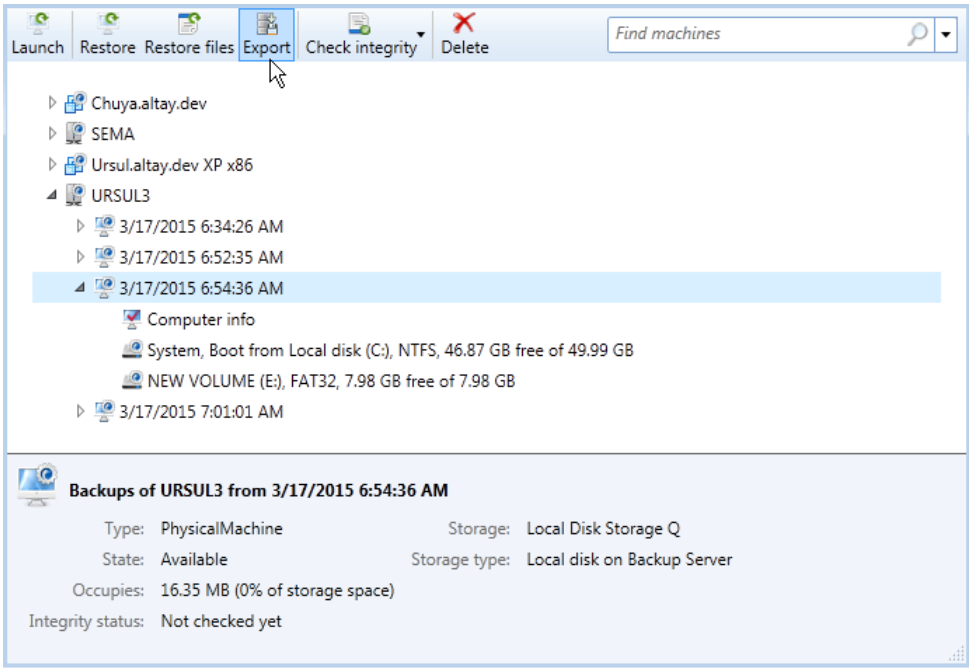

3. First the wizard will prompt you to select a machine from the list of all protected machines of all storages registered in the infrastructure.

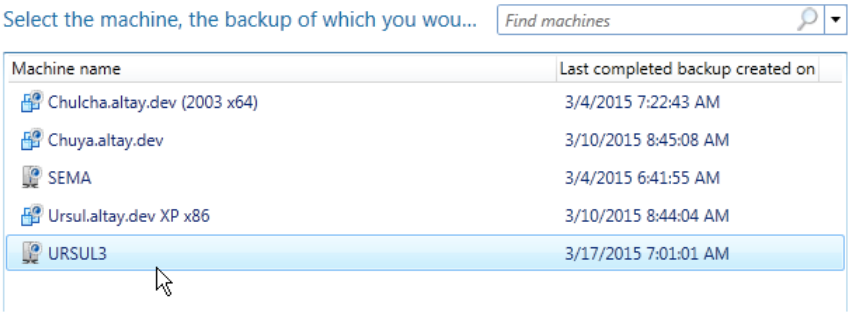

4. Then you will need to choose a desired restore point, if several available. If there are too many items on the list, filter the list through the 'Find snapshots created on' option by providing the exact date and time. By default, restore points from primary storages only will be available to work with. To view all restore points of the specified machine, please unmark the corresponding option.

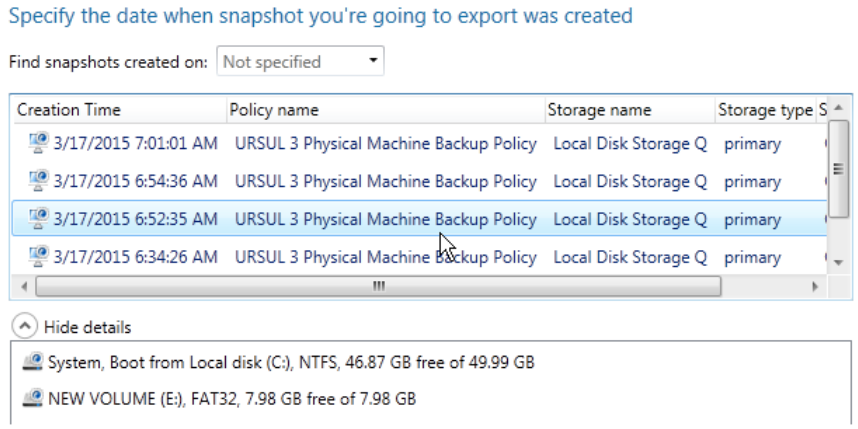

Steps 3 and 4 will only be available if you start the operation from the Backup & Restore ribbon.

5. Set up the export operation, and then click **Finish** to initiate.

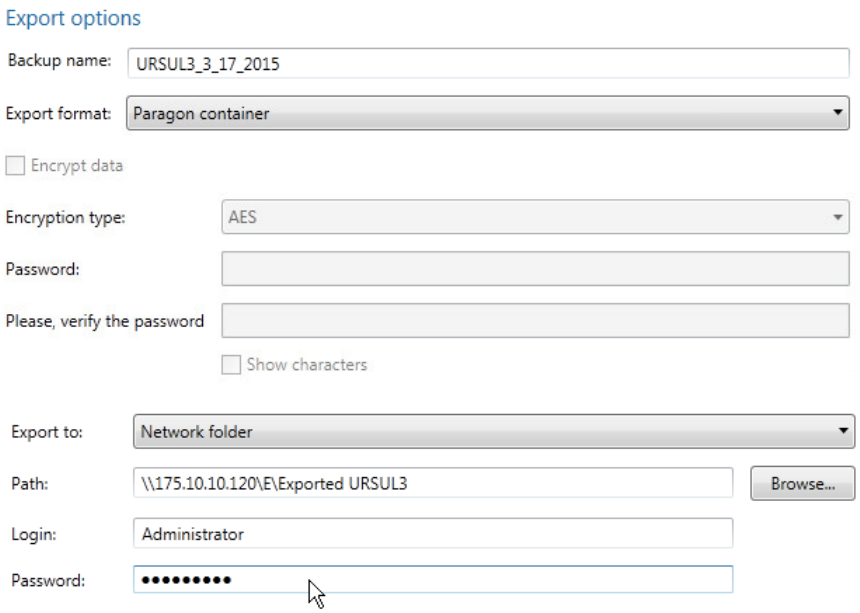

- **Backup name**. By default, the wizard offers a general name containing the name of the protected machine and the backup date, which you can modify however.
- **Encrypt data**. Use this option to protect the restore point against unauthorized access through an industry-standard 256-bit AES encryption algorithm. Please note that encryption is only available for Paragon's containers.
- **Export format**. The selected restore point can be exported in one of the supported formats (Paragon's container aka pVHD, VMWare VMDK, VHD, VHDX, or VirtualBox VDI). If selecting VHD or VHDX, you will receive not only a set of corresponding virtual containers, but a 90% ready-to-go virtual machine containing a configuration file, all necessary DLL libraries and Paragon's **Prm.PrepareForImport.exe** utility.

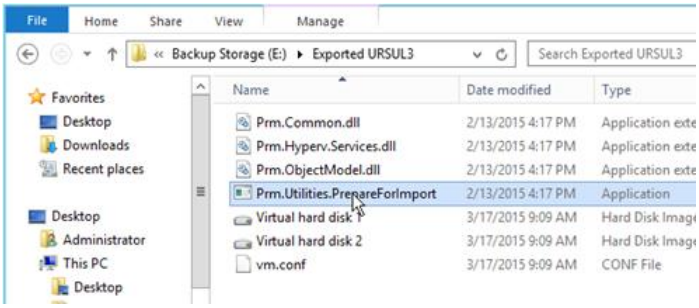

  For VHD it will be a Hyper-V V1 machine supported by MS Hyper-V Server 2008 R1/R2 and 2012 R1/R2, while for VHDX – Hyper-V V2 supported by MS Hyper-V Server 2012 R1/R2 only. Please consult the Launching Exported VMs on Hyper-V for more information.

- **Export to**. You can select either a local (for Backup Server) or network folder as destination of the exported backup data.
- **Path**. Click **Browse…** to specify exact location of the exported backup data.

6. The operation will be initiated immediately, which you can see through a popup window.

7. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.

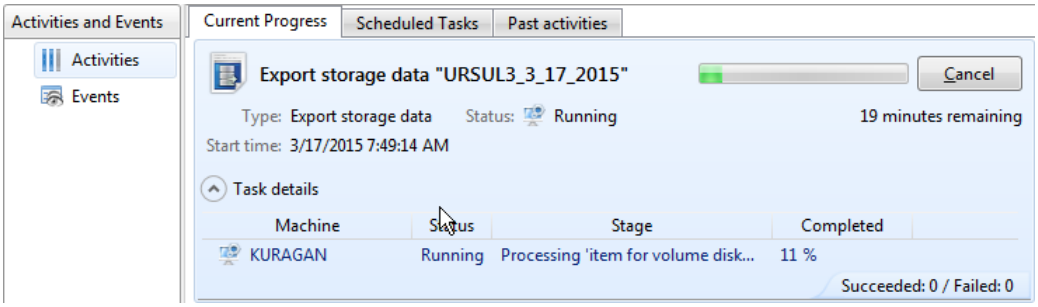

Launching Exported VMs on Hyper-V

1. Windows OS inside the resulted virtual machine is adjusted to start up on Hyper-V Server during the export procedure. However, once the exported virtual machine folder is on the hypervisor, you need to run **Prm.PrepareForImport.exe** there to create a Hyper-V VM (.xml) configuration file and a VM folder moving to corresponding subfolders all necessary VM files according to the used Hyper-V version.

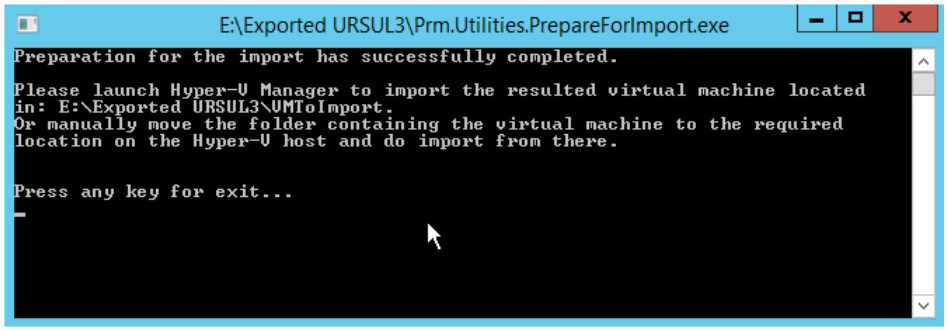

All activities of the command-line utility are saved to the 'result.log' file.

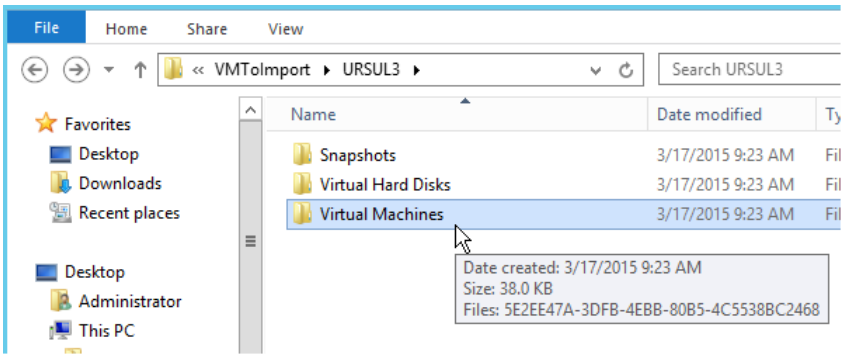

2. Once done, you can open the Hyper-V Manager and do a standard import VM operation specifying a folder named after the exported machine inside the **VMToImport** folder.

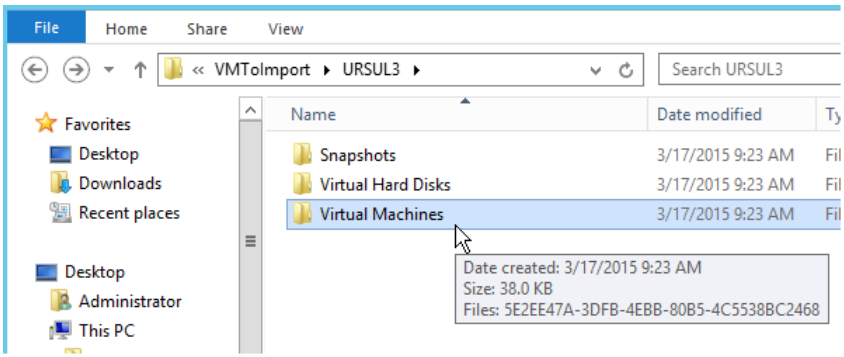

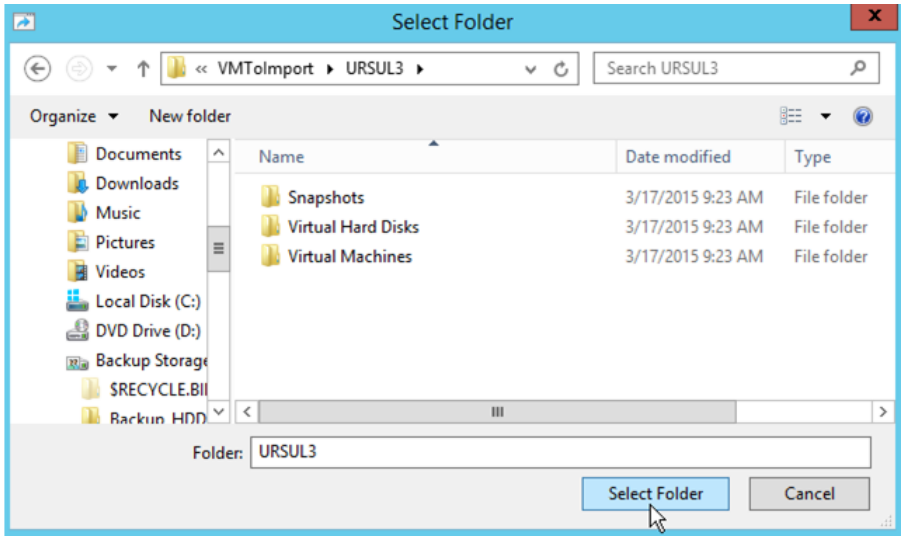

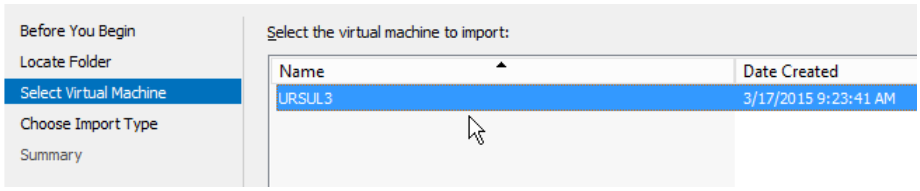

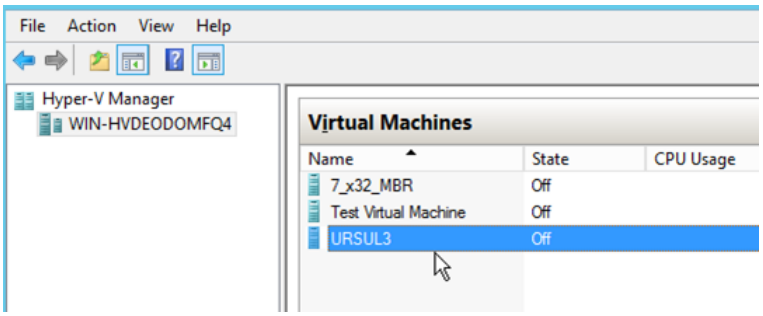

Managing Deduplication Server

About Paragon's Deduplication

Why Paragon's deduplication

Paragon's deduplication as any other deduplication mechanism aims at performing two basic tasks: it **detects** and **shares** duplicate data blocks instead of storing multiple copies. Different solutions do that differently and obviously with different deduplication efficiency, performance, and overheads. Paragon's deduplication is one of the most efficient on the market at the moment. Let's see why:

- Files are being split to unique blocks directly on protected machines (so called inline deduplication, which is done by the physical backup agent) or secondary backup storages (done by Backup Server in dual backup scenarios), which

enables to considerably cut on network traffic, while improving RPO. Most competitive solutions send full backup data to backup storage first, then start deduplication;

- In combination with Paragon's dual backup mechanism, backup data deduplication becomes efficient for near-CDP scenarios (VM backup through VMware CBT technology) when the inline deduplication is either impossible or considerably slows down performance of incremental imaging;
- Paragon's deduplication is global, i.e. all backup data (generated by various protection policies) sent to backup storages that are linked with one Deduplication Server is deduplicated. Competitive solutions offer a block-level deduplication for backup images created by one backup job;
- Deduplicated data is stored independently from backup storages. Location of Deduplication Server against Backup Server doesn't influence performance and deduplication efficiency during backup/dual backup or restore from deduplicated restore points;
- Paragon's deduplication works the same way for both, virtual and physical backup images, thus enabling to achieve better deduplication efficiency.

Inline deduplication during backup employs a different data reading mechanism, which is slower than the standard non-dedup backup. However, dedup backup goes much faster when deduplication storage contains a lot of unique blocks, as in this case slower reading is compensated by sending less data to the storage. Combining incremental imaging with deduplication is even more effective. But the best inline deduplication practice is backup through WAN links.

Paragon's deduplication peculiarities

Please take into account the following peculiarities of Paragon's deduplication mechanism:

- Only primary and secondary local and/or network backup storages can be linked to Deduplication Server;
- Only new images appeared on linked backup storages will be deduplicated;
- All types of images (physical or virtual) stored on secondary backup storages will be deduplicated;
- Backup objects containing unsupported file systems cannot be deduplicated;
- Files of less than 4KB in size will not be deduplicated to avoid unreasonable overheads;
- Deduplicated backup images are useless without an operating Deduplication Server and its storage;
- Immediate virtualization scenarios are not available for deduplicated backup images, while restore at file level is quite possible;
- NTFS-encrypted files (EFS) are not deduplicated at all and remain on backup storage.

Recommended deduplication system environment

- Deduplication storage may contain data blocks used by multiple backup sessions and backup storages. Thus loss of one shared data block leads to loss of all backup sessions dependent on it. This is why a regular integrity checkup of deduplication blocks is extremely critical. We offer to check deduplication blocks continuously in the background. Click here to know more about validation of deduplication blocks.
- It's allowed and highly recommended to configure two Deduplication Servers, thus you can set one of them as mirror or establish a special configuration where both deduplication servers act as exchangeable nods. Click here to know more about configuration of a mirror Deduplication Server.

- Deduplication block containers (.dvhd files) can also be protected by any third party solution, e.g. copied to any secured media, including tape.
- You can get the most out of deduplication by linking secondary storages to Deduplication Server, this way you can deduplicate all types of images (physical or virtual). Thus we highly recommend you to use Paragon's deduplication in combination with dual backup configurations.
- If you're using network-based deduplication storage, then decent performance can only be achieved through an isolated network link between Deduplication Server and a file server where deduplication storage is resided (separate subnet and network cards).

Registering Deduplication Storages

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Deduplication Server** is installed.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Deduplication Servers**.
3. If you've got several deduplication servers, first select the required dedup server by clicking on its name, then click on the **Configure Deduplication Server** link.

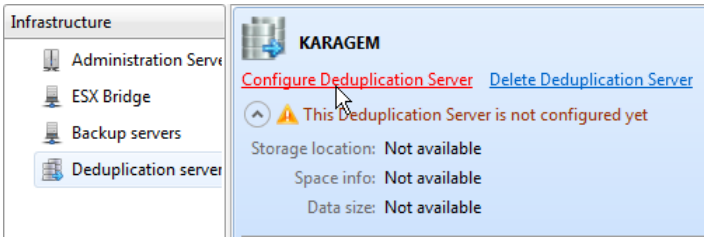

You can also initiate this operation by the right click of the mouse button, then selecting the corresponding option.

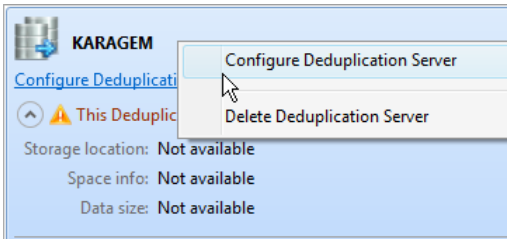

4. In the opened dialog you should first specify the required type of deduplication storage (**Local disk storage** or **Local network server**). Depending on your choice, you will be prompted to set a number of additional parameters:

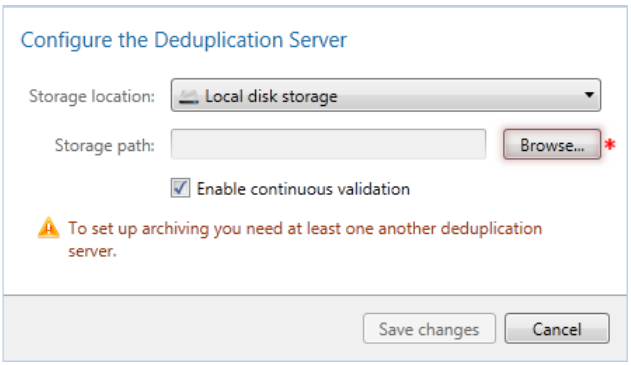

For local disk storages:

- **Storage path**. Click **Browse** to specify a local disk and folder on Deduplication Server to store deduplication blocks. Use the **New folder** button if necessary. Please make sure there's enough free space on the selected volume.

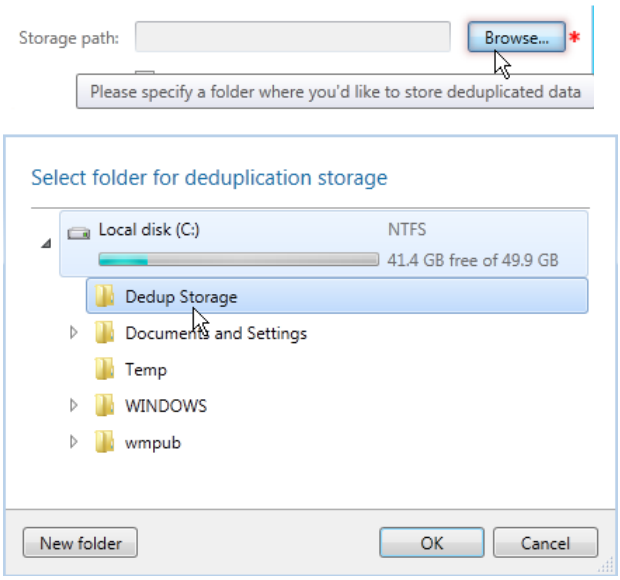

For local network storages:

- **Storage path**. Specify the required network share by manually entering its location or click **Browse** to find it on the net.

  Manually:

  Through browsing:

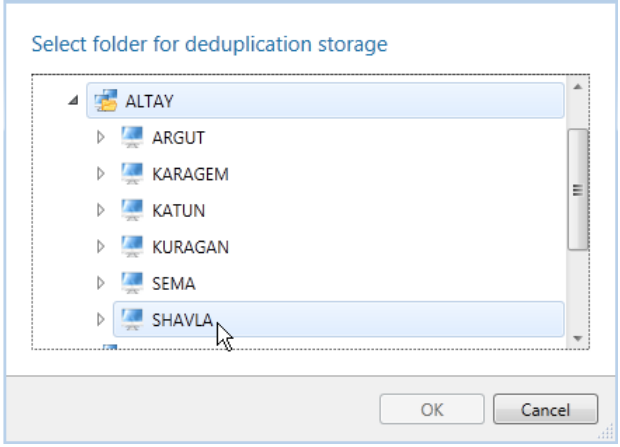

Double click on the required network machine to provide access credentials.

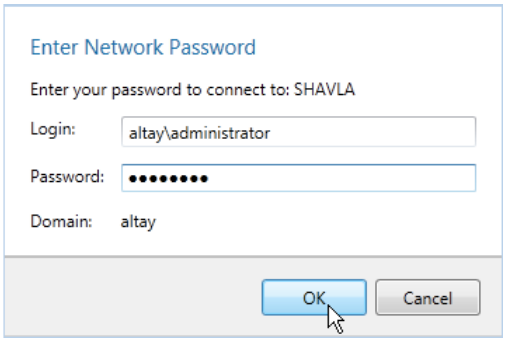

If the provided credentials are valid, you will be able to browse the specified network machine for the required storage folder. Click **OK** when ready.

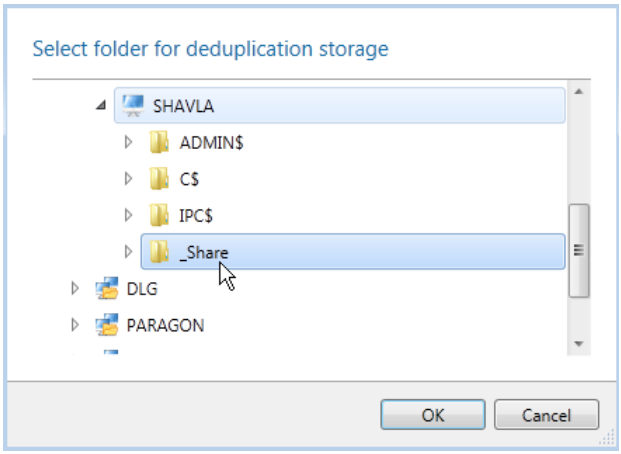

Decent performance can only be achieved through an isolated network link between Deduplication Server and a file server where deduplication storage is resided (separate subnet and network cards).

Due to certain limitation on simultaneous connections of a non-server OS, please make sure the network share specified as deduplication storage is hosted by a Windows Server machine.

- **Login and Password**. Specify access credentials for the manually provided network resource.

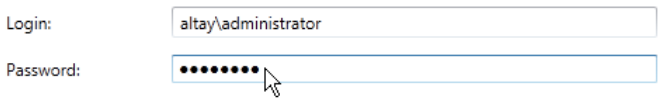

- **Enable continuous validation**. Unlike backup storage, deduplication storage may contain data blocks used by multiple backup sessions and backup storages. Thus loss of one shared data block leads to loss of all backup sessions dependent on it. This is why a regular integrity checkup of deduplication blocks is extremely critical. But it's a waste of time and hardware resources to check dedup blocks according to a certain backup session or machine, as in this case the same data blocks may be checked several times. To make this process effective, we offer to validate all dedup blocks of the storage in the cycle mode (default option), setting time stamps to each verified block. Through the **Trust interval** option (available in properties of a backup storage) the administrator can specify a time period during which blocks are assumed as consistent. Backup Server just analyzes these time stamps during a backup integrity checkup process to see whether this or that block can be trusted or not. If one of the questioned time stamps is too old, a forced checkup is initiated. This

continuous validation process goes in the background with lower priority than all other activities of Deduplication Server.

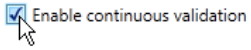

- **Enable archiving**. If you've got more than one deduplication server, there becomes available an additional option enabling to set one of the servers as mirror. We highly recommend you to have a mirrored dedup server configuration for better backup data protection. Deduplication blocks will start transferring to the mirror deduplication storage as soon as they are available on the source storage. Moreover, you've got the option to activate this option for both servers, thus establishing a special configuration where both deduplication servers act as exchangeable nods.

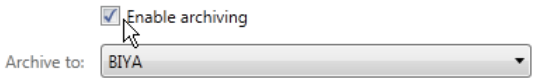

- Click **Save changes** to complete configuration of the deduplication storage.

5. As a result you should have a deduplication storage registered on the selected Deduplication Server.

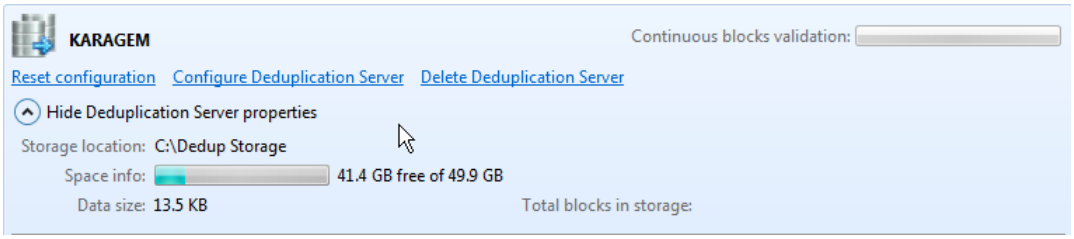

Linking Backup Storages to Deduplication Server

The last action you should take to allow backup data deduplication is to link required backup storages to Deduplication Server, this way enabling all newly created backup images on these backup storages to be automatically deduplicated.

Before you start deduplicating backup data, please consult the About Paragon's Deduplication chapter.

Once enabled, backup data deduplication cannot be turned off.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be registered at least one primary or secondary local/network backup storage.
- **Protect & Restore Deduplication Server** is installed and configured.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.
3. On the right pane select the required local or network storage (primary or secondary), then click the **Modify storage settings** icon or double click the storage.

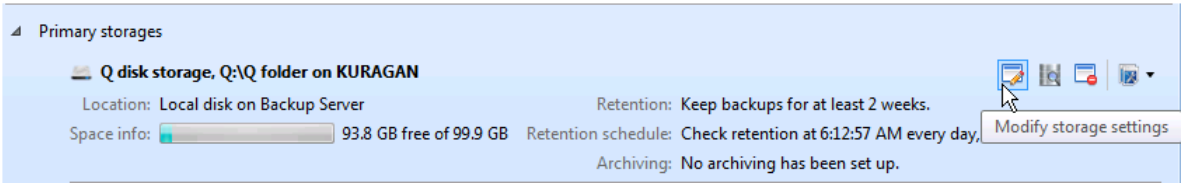

4. In the opened dialog click on the **Deduplication** tab to specify the following parameters:

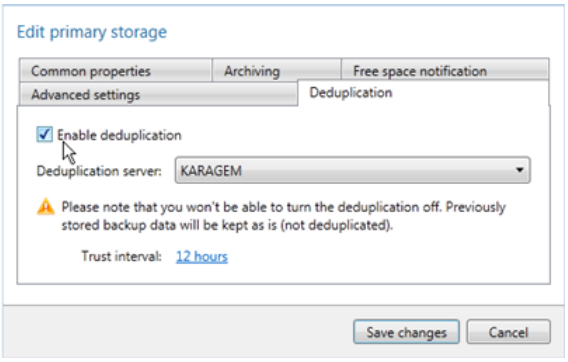

- **Enable deduplication**. Mark the checkbox to enable deduplication of backup data.
- **Deduplication server**. Select the required deduplication server (if several available).
- **Trust interval** (12 hours by default). Specify a time period during which deduplication blocks are assumed as consistent. The integrity checkup of a deduplicated backup image is organized in such a way that Backup Server analyzes time stamps of deduplication blocks that build up this image to see whether this or that block can be trusted or not. If one of the questioned time stamps is too old, a forced checkup is initiated. Click here to know more about validation of deduplication blocks.

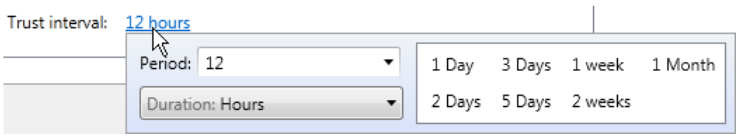

5. When ready, click **Save changes**. As a result you should have two additional properties for the specified backup storage (**Deduplication server** and **Deduplication efficiency**).

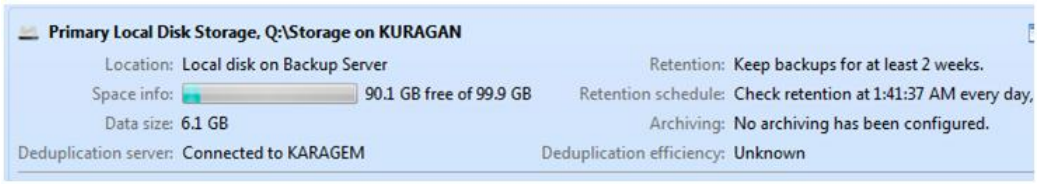

Understanding Deduplication Efficiency

PPR gives transparent results of deduplication efficiency, not taking into account backup image compression ratio. It's available for the user as percentage the original image size shrunk after deduplication. Deduplication efficiency is not only calculated for a single restore point, but for a backup catalog, or the entire backup storage.

You can see this information among other properties of a deduplicated (linked to Deduplication Server) backup storage (dedup efficiency of the entire storage):

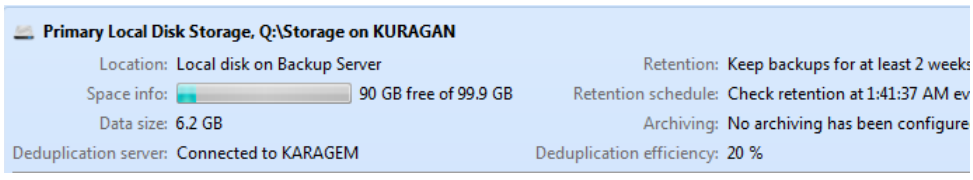

Or while browsing the storage and clicking on a certain machine (dedup efficiency of a single backup catalogue):

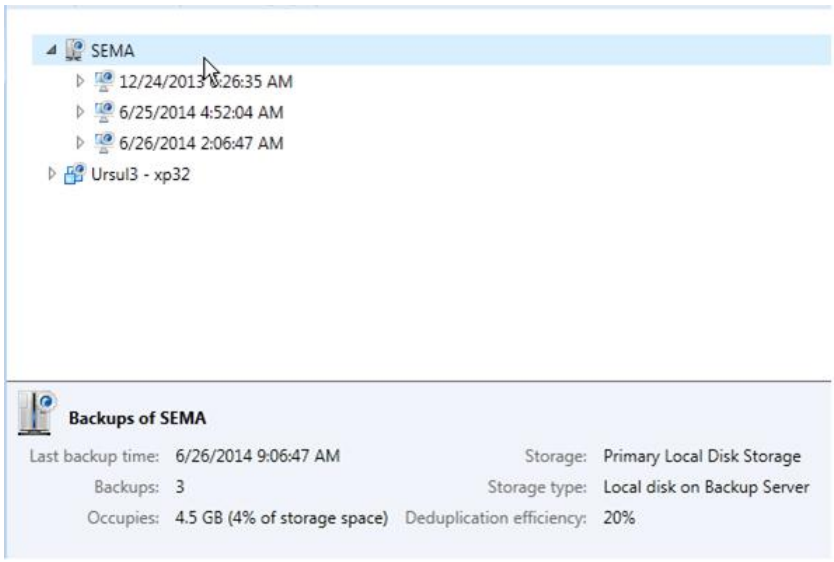

Or while browsing the storage and clicking on a certain backup image (dedup efficiency of a single restore point):

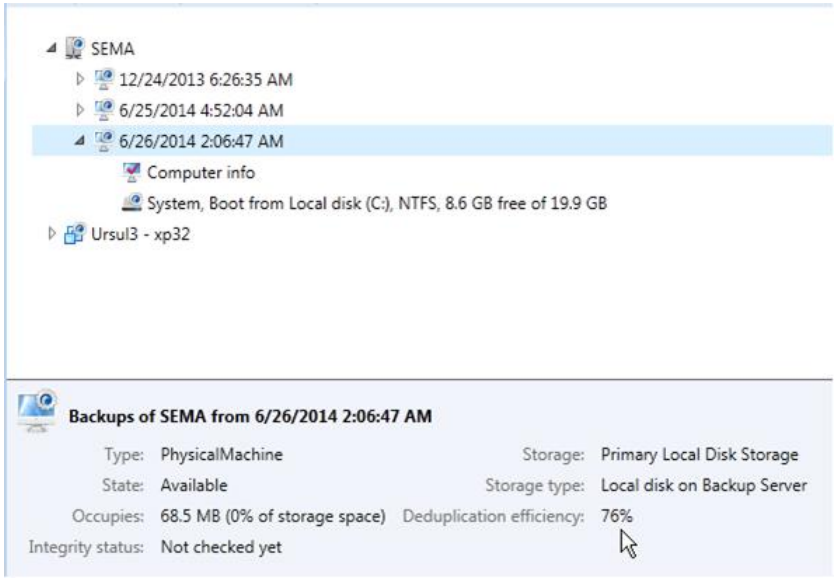

Obviously, deduplication efficiency increases with a number of deduplication blocks the dedup storage contains. However, if source machines contain too many files of less than 4KB or NTFS-encrypted files, deduplication efficiency will be low. For more information, please consult the About Paragon's Deduplication chapter.

Resetting Deduplication Storages

You can unregister (reset) any dedup storage from the infrastructure. This operation doesn't involve deletion of deduplication blocks of the specified storage, but only its removal from the infrastructure. Later you can register this storage again by pointing to the same location during configuration of a new deduplication storage. This option can help to easily migrate deduplication storage.

To reset an existing deduplication storage, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Deduplication Servers**.
3. On the right pane select the required storage, then click the **Reset configuration** link.

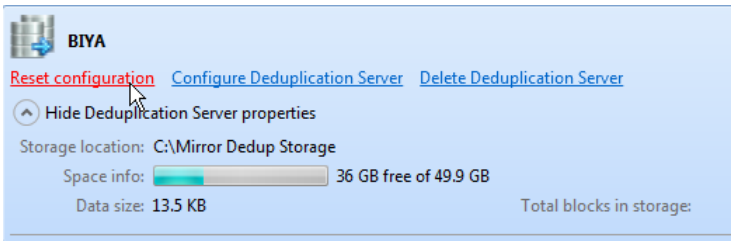

4. The operation will be initiated immediately.

Migrating Deduplication Storages

To migrate an existing deduplication storage to another location, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Deduplication Servers**.
3. Reset the required deduplication storage.
4. Move the folder that contains deduplication blocks to another location.
5. Register a new deduplication storage specifying the just moved folder in the **Storage path** field.
6. That's it. Deduplication Server will take in use deduplication blocks from the specified location.

Changing Deduplication Server

You're allowed to change Deduplication Server for existing linked backup storages that contain deduplicated restore points. This operation can help you easily swap to a mirror deduplication server or change a dedup server after dedup storage migration. However, you should understand pretty well possible aftereffects of this action. If you link a backup storage with deduplicated backup images to some other deduplication server, which deduplication storage doesn't contain necessary deduplication blocks, then these images will be detached from corresponding deduplication blocks, thus you won't be able to restore from them. If you'd like to change Deduplication Server, first please make sure the new server contains required dedup data.

If you want the new Deduplication Server to only have deduplicated data of particular restore points from the interested backup storage, please create a new backup storage, link this storage to the new dedup server, and then copy required restore points to it through the archiving functionality.

Prerequisites

- **Protect & Restore Backup Server** is installed.
- There should be registered at least one primary or secondary local/network backup storage.
- There are two Deduplication Servers in the infrastructure.
- Each Deduplication Server has registered deduplication storage.
- At least one backup storage should be linked to one of the Deduplication Servers and contain deduplicated backup images.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Backup Servers**.

3. On the right pane select the required local or network storage (primary or secondary), then click the **Modify storage settings** icon or double click the storage.

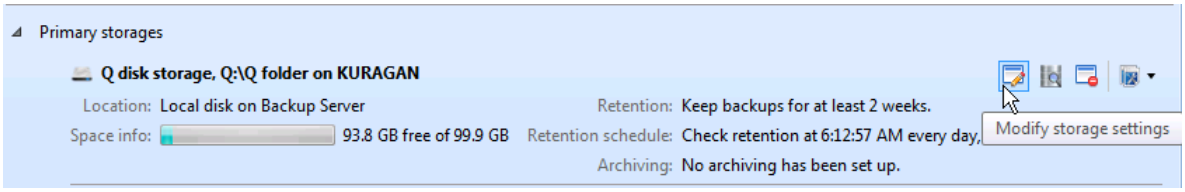

4. In the opened dialog click on the **Deduplication** tab to specify the following parameters:

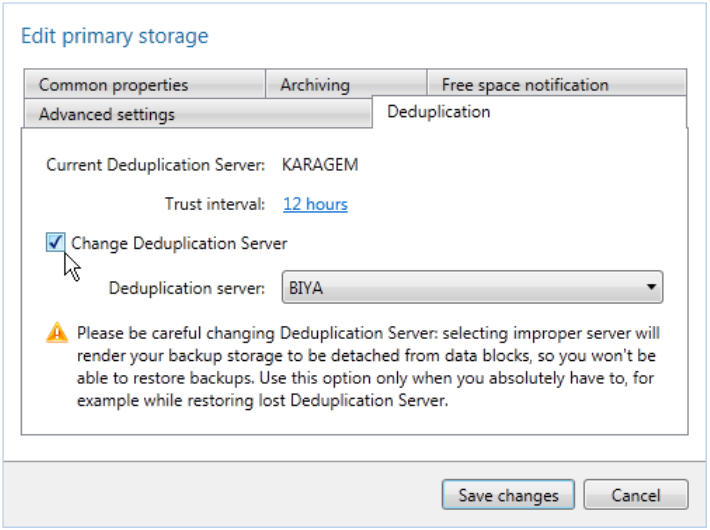

- **Change Deduplication Server**. Mark the checkbox to change the server.
- **Deduplication server**. Select the required deduplication server (if more than two are available).

If you link a backup storage with deduplicated backup images to some other Deduplication Server, which deduplication storage doesn't contain necessary deduplication blocks, then these images will be detached from corresponding deduplication blocks, thus you won't be able to restore from them.

---

5. When ready, click **Save changes**.

Deleting Deduplication Server

Beside removal of the Deduplication Server role from a certain machine, this operation involves deletion of deduplication blocks of the corresponding dedup storage.

To delete an existing deduplication server, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Infrastructure > Deduplication Servers**.
3. On the right pane select the required storage, then click the **Delete Deduplication Server** link.

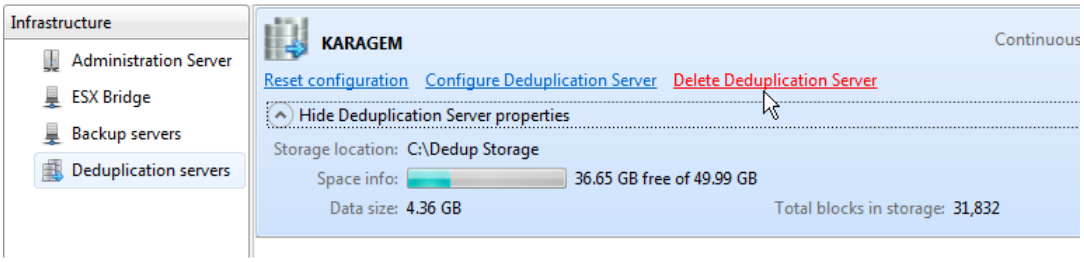

You can also initiate this operation by the right click of the mouse button, then selecting the corresponding option.

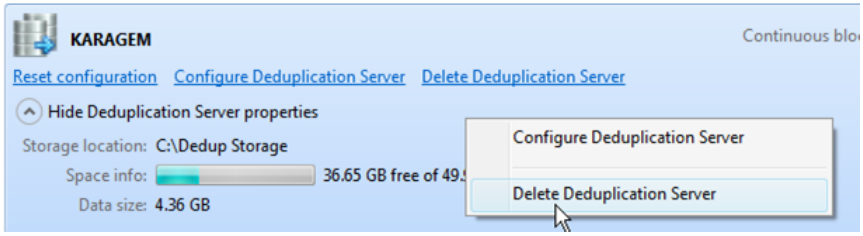

4. You will be informed on backup storages (if any) that are linked to this server and have deduplicated restore points, which you will be unable to restore if proceeding with the deletion. If you need to completely remove this machine from the infrastructure with all existing roles it has, please additionally mark the corresponding option. Click **Delete** to confirm the operation.

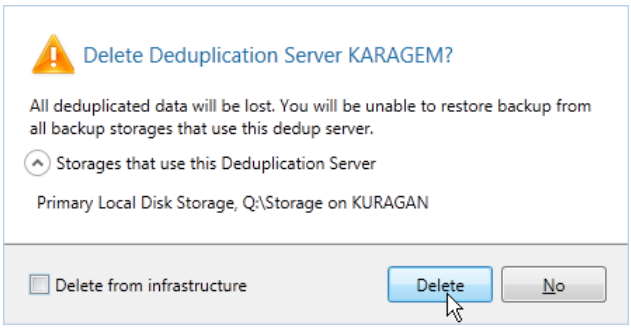

Protecting Virtual Machines

Backing up Virtual Machines

PPR enables to back up any Windows, Linux or other OS guest supported by VMware. One backup task can involve one or many virtual machines. By default, for every machine our product creates a full backup in a special proprietary format during the first run, then incremental updates according to a set timetable. It allows configuring general retention policies for backup storages or a particular policy for a certain backup task, specifying how long backups should be kept or the amount of space they can take. When time comes, all restore points beyond the set limit are merged with their full backup thus creating a new full backup. All backup images are being highly compressed during creation by using redundant data exclusion filters (OS page files, zero data blocks, etc.), which eases the backup storage requirements.

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore ESX Agent** is installed in a guest environment of an ESX server that hosts virtual machines you're planning to protect. Besides at least one ESX host is registered on it.
- **Protect & Restore Backup Server** is installed on any machine, but the more powerful, the better.

- There has been registered a primary local or network backup storage. To know more on the subject, please consult the Registering primary storages chapter.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **New VM Backup Policy**,

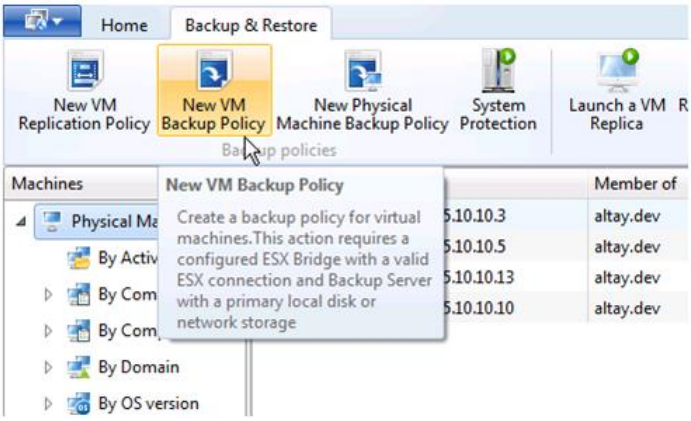

or go to **Policies >** right click on the **Backup policies for virtual machines**, then select **New VM Backup Policy**.

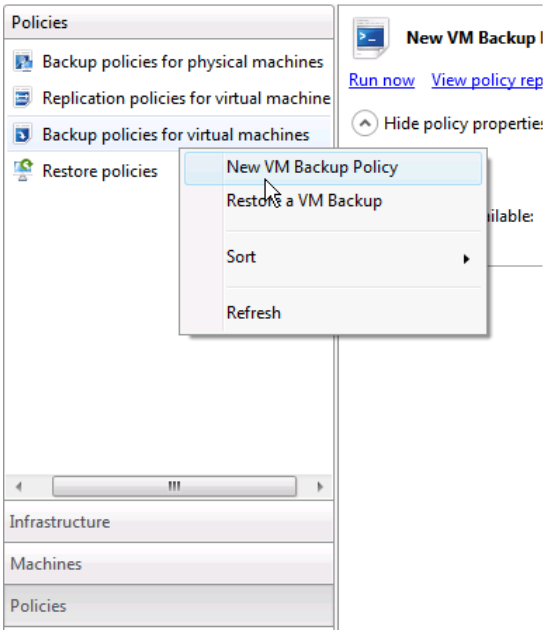

3. The opened dialog consists of two tabs that include a number of parameters:

   The first tab (Policy details):

   - **Policy name**. Give it a catchy name.

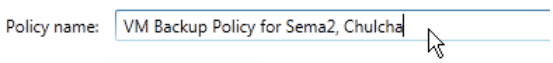

   - **Description**. Give a detailed description to the backup task (optional).

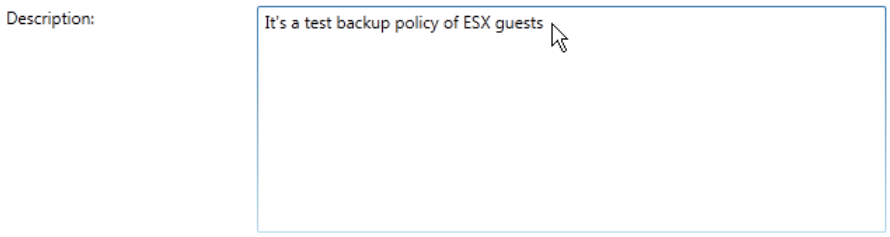

- **Back up to**. Select a backup server (if several), then the required primary storage from the popup list to place backup images to.

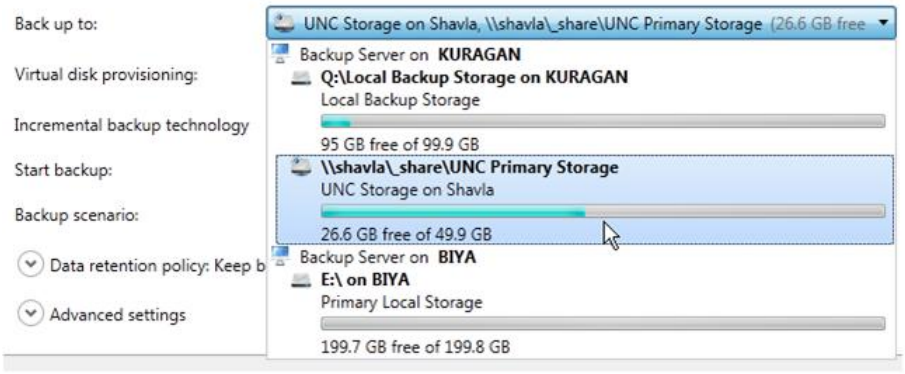

- **Virtual disk provisioning**. Specify a disk provisioning type. VMware ESX/ESXi hosts support two types of disk provisioning, namely “Thin” and “Thick”. With the “thin provisioning” the size of a VDMK file on a datastore is exactly the amount of data it contains, so if you create a 300GB virtual disk, and place 50GB of data in it, the VMDK file will be 50GB in size. With the “thick provisioning” the size of a VMDK file on a datastore is always its maximum size, so no matter how much data the virtual disk contains, the VMDK file will be 300GB in size anyway. Obviously, conversion from the “thick provisioning” to the “thin” when replicating disks of virtual machines may significantly cut backup storage requirements and costs.

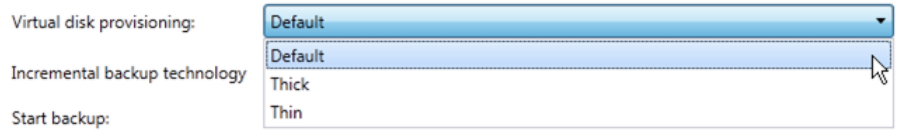

- **Incremental backup technology**. By default, for every machine our product creates a full backup for the first run, and then only saves changes since the last performed operation in incremental images. The delta to write is either parsed through VMware CBT or Paragon’s ITE. So here you’ve got five options to choose from:

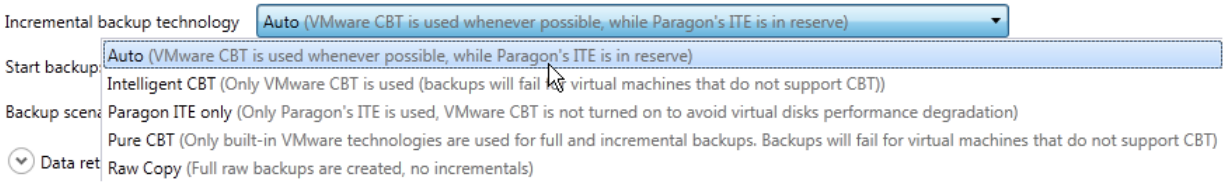

- **Auto**. It’s the default mode, when **Intelligent CBT** is used whenever possible, while **Paragon’s ITE** is in reserve. Please note that in situations when CBT is turned off on the target machine, but this machine contains file systems unsupported by Paragon, the **Raw Copy** mode is automatically activated, which doesn’t allow incremental imaging at all.
- **Intelligent CBT**. In this mode changes since the last backup are parsed through VMware CBT, and then this data is considerably reduced through Paragon’s patent-pending algorithms, thus producing a much

smaller backup image. If VMware CBT is turned off on the target machine, backup tasks will fail with a corresponding warning.

- **Paragon ITE**. In this mode changes since the last backup are parsed through Paragon's ITE only. Resulted backup images will take a bit more time to create and be larger in size, but it's the only decent option when CBT cannot be used on the target machine. Another benefit comes from the fact that an active CBT significantly degrades the disk subsystem performance of target machines.
- **Pure CBT**. In this mode changes since the last backup are parsed through VMware CBT only and then saved in the resulted image without any optimization.
- **Raw Copy**. Use this mode if none of the other options can help you back up target machines. Only full images will be created.

- **Start backup**. By default, no schedule is set for the backup policy, so you will need to manually commit it after its validation. If you want to schedule the policy, just click on the corresponding link to specify a timetable.

The opened dialog consists of two sections:

Basic scheduling

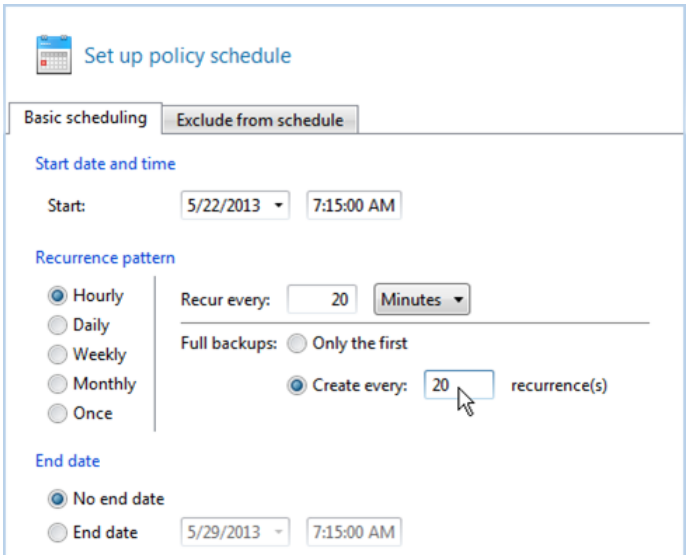

In this section you can set up a backup timetable. By default, a full backup will be created once for every target machine, then only come incremental updates, which you can change however through the **Full backups** section.

Exclude from schedule

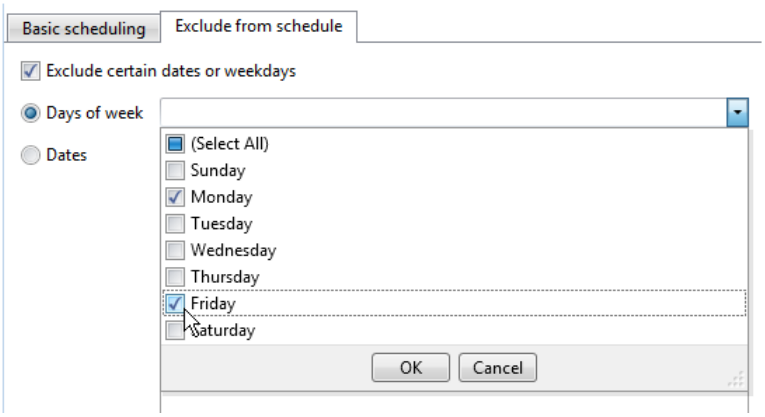

In this section you can specify days of week, or certain dates, when backup operations should not be accomplished.

If you schedule a backup task, the operation will start according to time of ESX Agent.

- **Backup scenario**. In the current version of the product only one backup scenario is supported (simple).

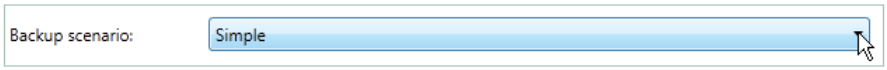

- **Data retention options**. Here you can specify a custom backup data retention mode that will be taken into account for the created policy only.

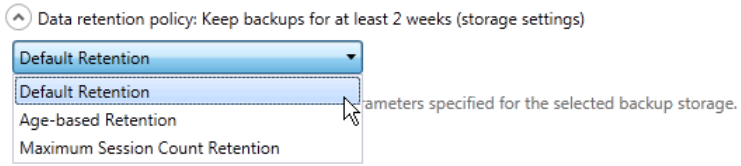

  - **Default Retention** to use data retention parameters specified for the selected backup storage.
  - **Age-based Retention.** Use this option if you'd like to limit lifetime of backup images created by this policy (2 weeks by default, set in the **Age-based criterion** option). To minimize load on the backup server, there's a conditional criterion (**Size-based criterion**) you can make use of to suppress the data retention process until size of backups per each machine exceeds a certain value (10 GB by default).

  - **Maximum Session Count Retention** to define the maximum number of backup sessions allowed for target machines processed by this policy. On exceeding the set value, backup chains will be thinned out starting from the oldest backup images.

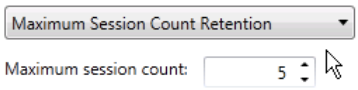

- **Advanced settings**. Click on the corresponding arrow button to see and configure additional options if necessary:

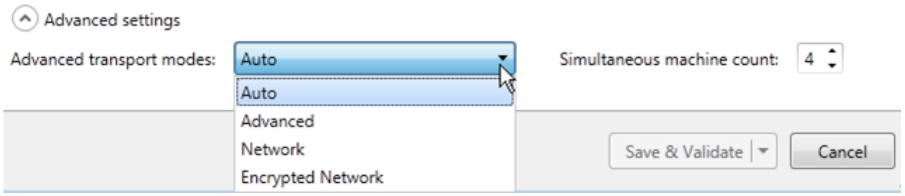

- ***Advanced transport modes***. There are several modes VMware vStorage API can provide access to data inside disk containers of virtual machines to do backup, restore, or partition alignment. Select the required one from the popup list (**Auto** by default):
    - ***Auto***. VMware picks the fastest and most efficient mode by trying each one on the fly (recommended);
    - ***Advanced***. It's also an auto mode, but here only two high-performance modes are considered to use (Virtual Appliance or SAN). **Virtual Appliance** will be selected if ESX Agent is deployed on a virtual machine hosted by ESX, which guests are being protected or optimized. **SAN** will be selected if ESX Agent is deployed on a physical machine running Windows Server 2003 / 2008 / 2012, which have access to LUNs (Logical Unit Number) with virtual machines. Both modes do not load CPU of ESX. If Virtual Appliance is configured not to use the production networking, then the production bandwidth is also not used. However, both these modes require special configuration of virtual machines.
    - ***Network***. It's a universal mode that can be used no matter where ESX Agent is deployed. It supports all configurations of virtual machines and can process virtual disks that do not support the snapshot mechanism (Independent or Physical RDM disks). Please note however that this mode uses the external network interface of ESX, thus it can heavily load its CPU and bandwidth. Besides, VMware limits for it the number of simultaneously processed disks by all parallel vStrorage API activities (32 disks for ESX 5.x, 27 for ESX 4.x, 23 for ESXi 4.x).
    - ***Encrypted Network***. It's the network mode additionally encrypted through SSL (Secure Sockets Layer). Obviously it's the slowest mode of all. Besides it requires special configuration of virtual machines. VMware doesn't limit the number of simultaneously processed virtual disks for it. We recommend this mode if ESX Agent connects ESX via public network with low level of confidence.
- ***Simultaneous machine count***. By default, only four virtual machines are allowed to back up simultaneously, which you can change however. Please note the maximum available value is 20 machines.

Advanced settings will be available if the corresponding option is enabled in the Settings dialog.

The second tab (VMs to back up):

- Browse the connected ESX to mark virtual machines you're going to protect.

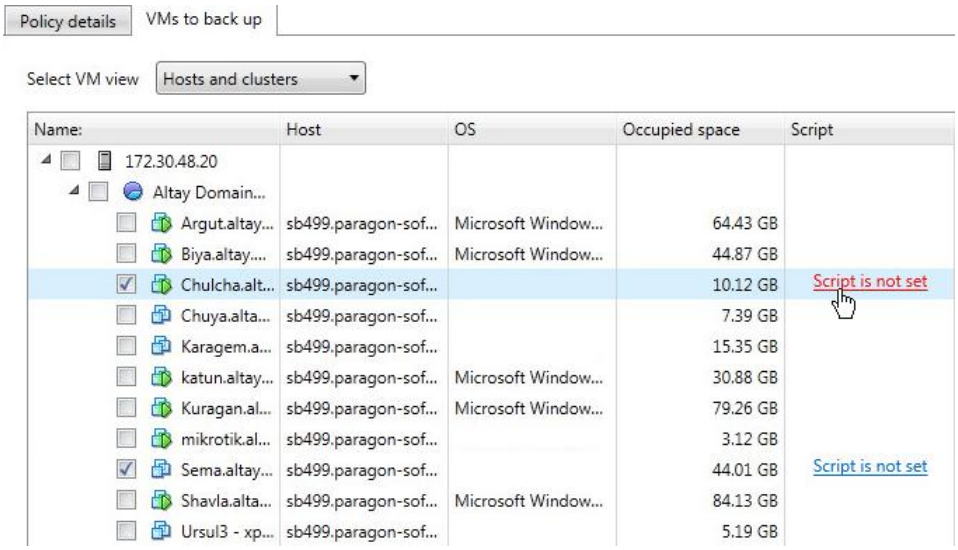

If guest machines you're going to protect run applications that do not support Microsoft VSS (an old version of MS SQL Server, Linux-based PostgreSQL or Oracle Database, etc.), you need to run custom scripts to provide a coherent state of all open files and databases involved in a backup. Pre-scripts help to properly freeze (quiesce) applications before PPR initiates creation of a snapshot, while post-scripts bring these applications back to normal work.

Currently supported script formats:

For Windows VMs	For Linux VMs
.cmd	.bash, .sh, .tcsh (Shell scripts)
.bat	.php (Perl scripts)
.js (Java scripts)	
.vbs (Visual Basic scripts)	

To specify pre- and post-scripts for a given virtual machine, please do the following:

- Click the **Script is not set** hyperlink opposite a marked machine.
- Set credentials of a local user with enough privileges to allow PPR log in to the machine and run scripts.

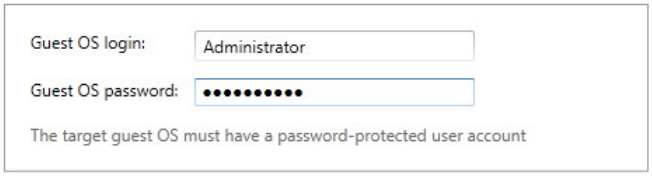

- Select **Enable script execution** to specify a pre-script. You can either set an absolute path to the early prepared script file that is stored on the target machine or enter commands directly in the corresponding textual field. Repeat the same actions for a post-script by clicking the **Post-script** tab.

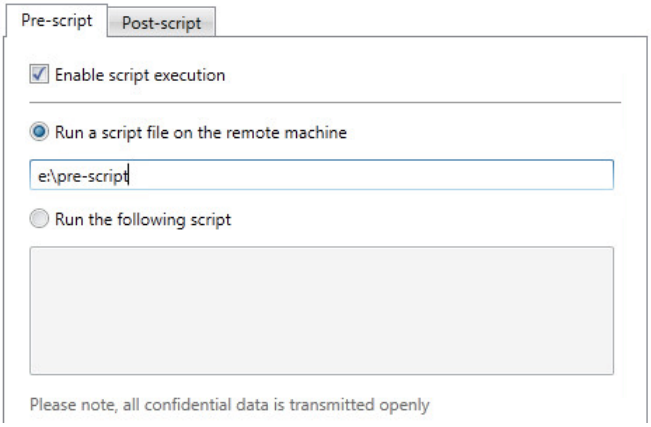

- Click **More options** to change default settings if necessary. A script is considered to be executed successfully if "0" is returned, which you can change to any value. If your scripts can't be processed by OS directly, set a path to the corresponding scripting engine. Besides you can change default retry and timeout intervals. By clicking **Delete on success** you allow deletion of scripts once they are executed.

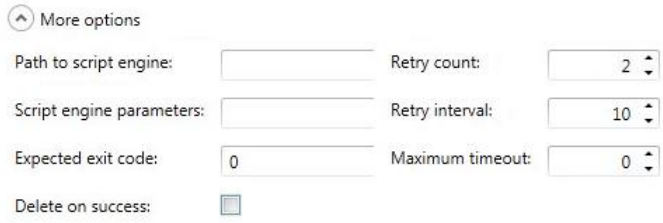

When you're ready with all parameters, click **Save & Validate** to complete creation of the backup policy. By default there will be used the fast level of validation, which you can change by clicking on the arrow button.

Let's see how three validation levels differ:

- **Fast**. It includes checkup of all policy rules and their parameters, availability of the backup storage and ESX connection parameters.
- **Medium**. It includes connection to the specified ESX host to scan for target virtual machines as well as connection to the backup storage to retrieve metadata from it.
- **Thorough and slow**. It includes creation/deletion of snapshots of target virtual machines, creation of an uncompleted backup session and data items in the backup storage without opening data streams and data copying.

Pre- and post-scripts are also processed during validation launches.

4. Validation of the backup task will be initiated immediately. You will be informed on the operation start through a popup window.

5. Additionally we recommend you to create a notification sub-policy for easier monitoring.

6. If the just created policy hasn't been scheduled, you need to manually commit it once the validation is over. To do that, please go to **Policies > Backup policies for virtual machines**, then select **Run now** for the corresponding policy.

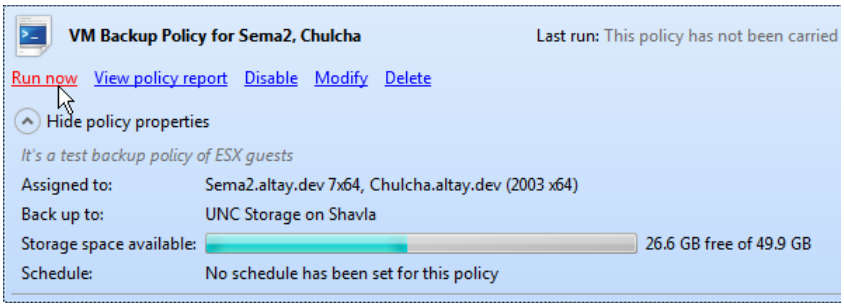

If it's not the first time you commit a backup policy, you will be offered to choose the required backup mode.

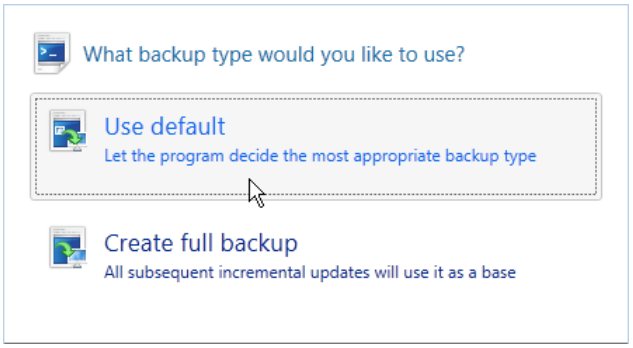

7. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.

8. When the backup task is over, its status will be updated.

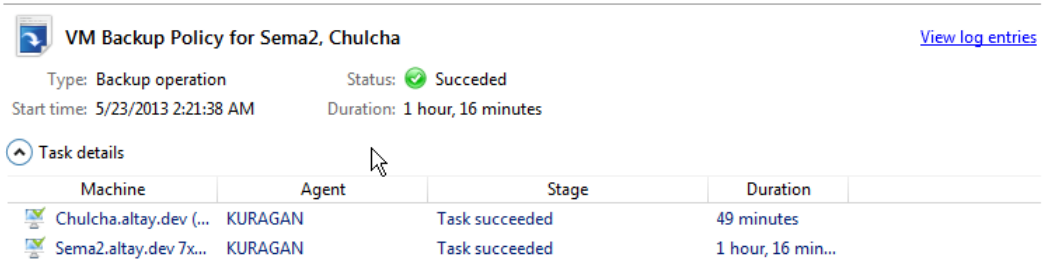

9. If going to the target backup storage now, you can see two new folders with guid names that contain backup data of the corresponding target virtual machines.

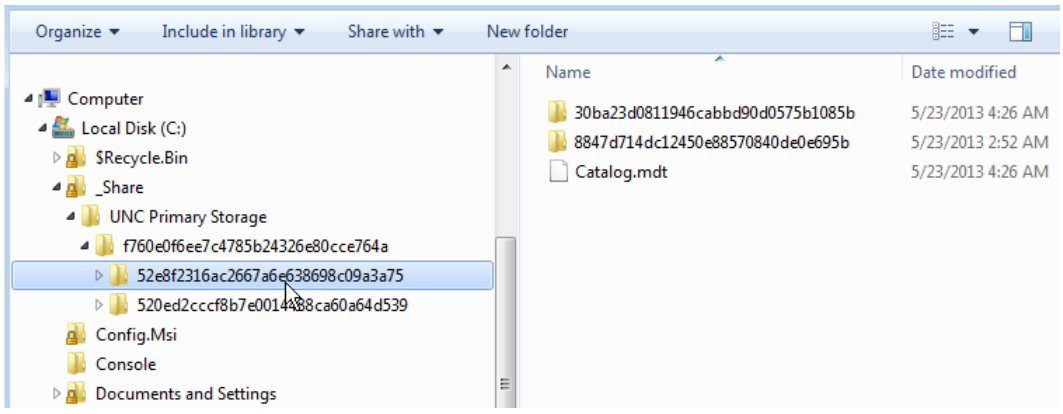

To avoid malfunctioning of the infrastructure, please administer backup data only through the PPR interfaces.

To know how to manage created policies, please consult the Managing Policies chapter.

Creating Replicas of Virtual Machines

For high-availability virtual environments that run the first tier applications, PPR complements VM Backup with VM Replication. Replication provides the best RTO (Recovery Time Objective), for this technique implies creation of clones (replicas) of target machines on a certain ESX datastore and their registering on the host under different names. Replicas are stored uncompressed in their native format, thus they are ready-to-go at any moment. All changes since the initial full replica are written to VMware native snapshot files, acting as restore points, thus allowing the usage of the VMware revert-to-snapshot mechanism to further accelerate disaster recovery scenarios, providing for almost zero downtime operation. You can also define a retention policy for replicas, thus all snapshots that breach the set policy will be automatically collapsed.

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore ESX Agent** is installed in a guest environment of an ESX server that hosts virtual machines you're planning to protect. Besides at least one ESX host is registered on it.
- **Protect & Restore Backup Server** is installed on any machine, but the more powerful, the better.
- There has been registered a primary ESX storage. To know more on the subject, please consult the Registering primary storages chapter.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **New VM Replication Policy**,

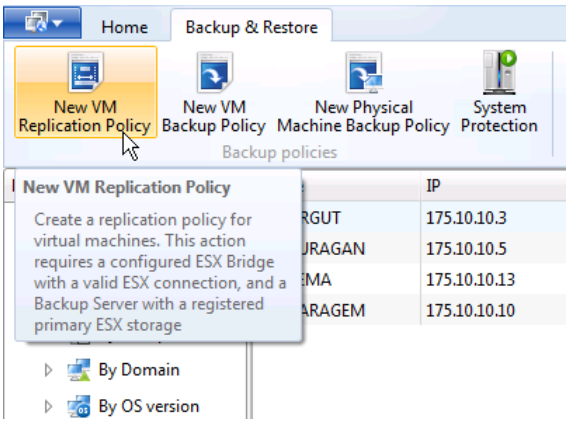

or go to **Policies >** right click on the **Replication policies for virtual machines**, then select **New VM Replication Policy**.

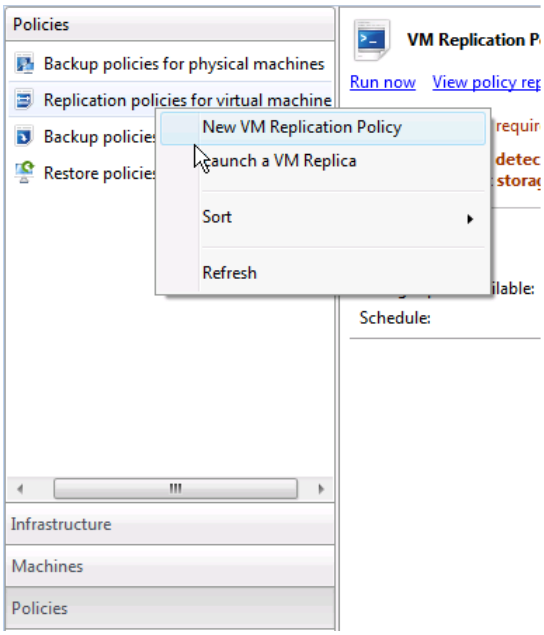

3. The opened dialog consists of two tabs that include a number of parameters:

   The first tab (How to replicate):

   - **Policy name**. Give it a catchy name.

     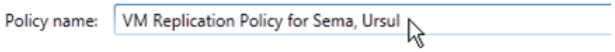

   - **Description**. Give a detailed description to the replication task (optional).

     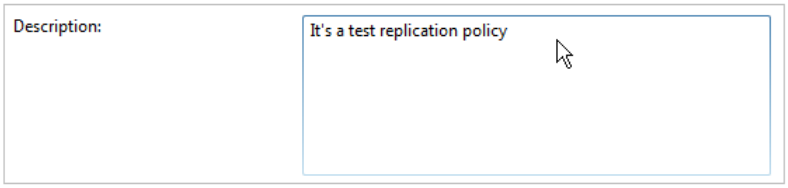

   - **Replicate to**. Select a backup server (if several), then the required primary ESX storage from the popup list to place replicas to.

     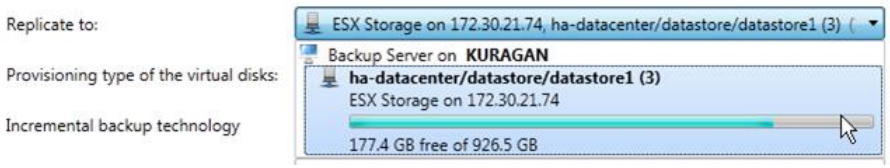

   - **Virtual disk provisioning**. Specify a disk provisioning type. VMware ESX/ESXi hosts support two types of disk provisioning, namely “Thin” and “Thick”. With the “thin provisioning” the size of a VDMK file on a datastore is exactly the amount of data it contains, so if you create a 300GB virtual disk, and place 50GB of data in it, the VMDK file will be 50GB in size. With the “thick provisioning” the size of a VMDK file on a datastore is always its maximum size, so no matter how much data the virtual disk contains, the VMDK file will be 300GB in size anyway. Obviously, conversion from the “thick provisioning” to the “thin” when backing up disks of virtual machines may significantly cut backup storage requirements and costs.

     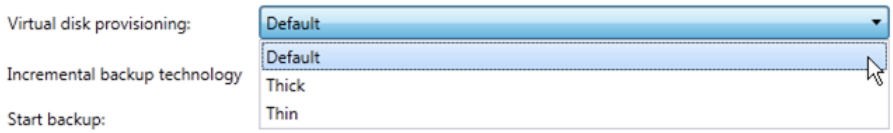

- **Incremental backup technology**. By default, for every machine our product creates a full replica for the first run, and then only saves changes since the last performed operation in incremental images. The delta to write is either parsed through VMware CBT or Paragon's ITE. So here you've got five options to choose from:

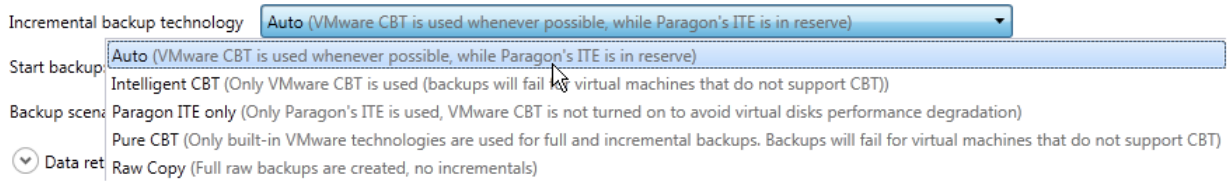

  - **Auto**. It's the default mode, when **Intelligent CBT** is used whenever possible, while **Paragon's ITE** is in reserve. Please note that in situations when CBT is turned off on the target machine, but this machine contains file systems unsupported by Paragon, the **Raw Copy** mode is automatically activated, which doesn't allow incremental imaging at all.
  - **Intelligent CBT**. In this mode changes since the last backup are parsed through VMware CBT, and then this data is considerably reduced through Paragon's patent-pending algorithms, thus producing a much smaller backup image. If VMware CBT is turned off on the target machine, backup tasks will fail with a corresponding warning.
  - **Paragon ITE**. In this mode changes since the last backup are parsed through Paragon's ITE only. Resulted backup images will take a bit more time to create and be larger in size, but it's the only decent option when CBT cannot be used on the target machine. Another benefit comes from the fact that an active CBT significantly degrades the disk subsystem performance of target machines.
  - **Pure CBT**. In this mode changes since the last backup are parsed through VMware CBT only and then saved in the resulted image without any optimization.
  - **Raw Copy**. Use this mode if none of the other options can help you back up target machines. Only full images will be created.

- **Replica suffix**. By default the suffix "_replica" will be added to the names of all replica virtual machines, which you can change here.

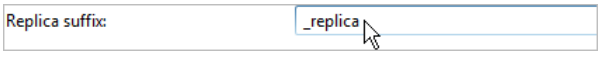

- **Start replica**. By default, no schedule is set for the replication policy, so you will need to manually commit it after its validation. If you want to schedule the policy, just click on the corresponding link to specify a timetable.

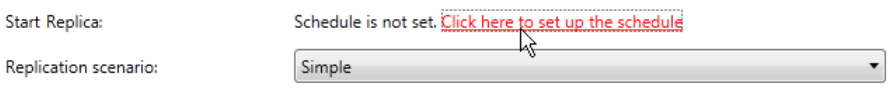

The opened dialog consists of two sections:

Basic scheduling

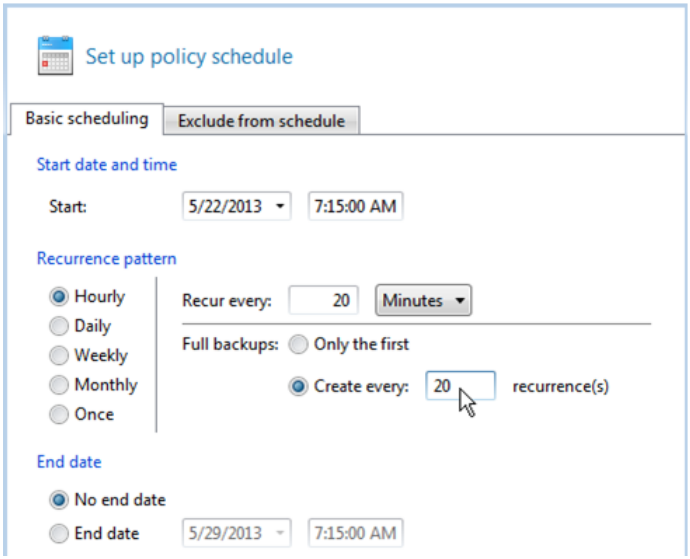

In this section you can set up a replication timetable. The minimal available update interval is one minute. By default, a full replica will be created once for every target machine, then only come incremental updates, which you can change however through the **Full backups** section.

Exclude from schedule

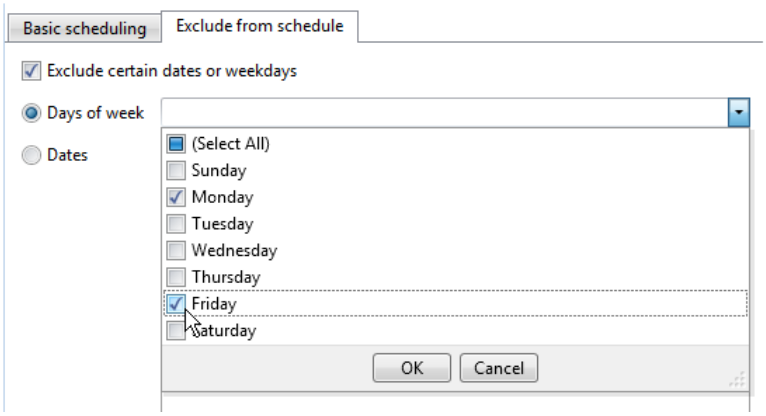

In this section you can specify days of week, or certain dates, when replica operations should not be accomplished.

If you schedule a replication task, the operation will start according to time of ESX Agent.

- **Replication scenario**. In the current version of the product only one replication scenario is supported (simple).

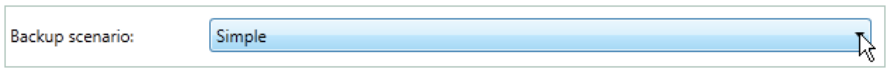

- **Data retention options**. Here you can specify a custom replica data retention mode that will be taken into account for the created policy only.

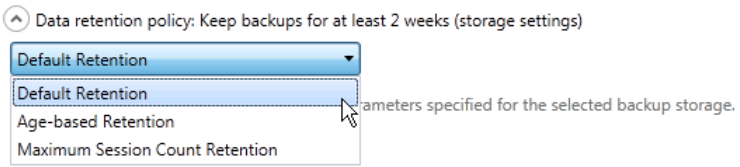

- **Default Retention** to use data retention parameters specified for the selected replica storage.
- **Age-based Retention.** Use this option if you'd like to limit lifetime of backup images created by this policy (2 weeks by default, set in the **Age-based criterion** option). To minimize load on the backup server, there's a conditional criterion (**Size-based criterion**) you can make use of to suppress the data retention process until size of backups per each machine exceeds a certain value (10 GB by default).

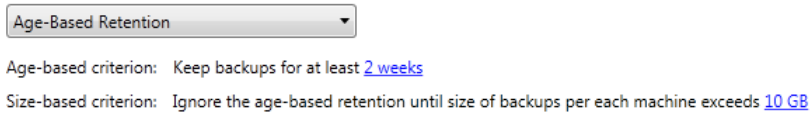

- **Maximum Session Count Retention** to define the maximum number of backup sessions allowed for target machines processed by this policy. On exceeding the set value, replica chains will be thinned out starting from the oldest replicas.

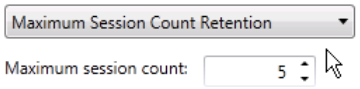

- **Advanced settings**. Click on the corresponding arrow button to see and configure additional options if necessary:

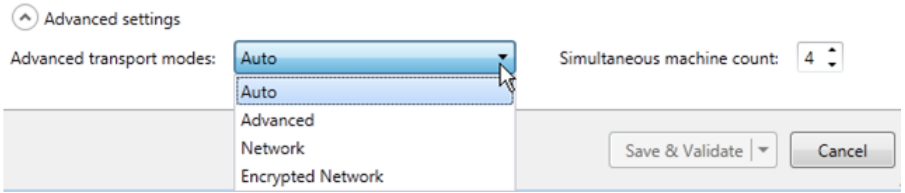

- ***Advanced transport modes***. There are several modes VMware vStorage API can provide access to data inside disk containers of virtual machines to do backup, restore, or partition alignment. Select the required one from the popup list (**Auto** by default):
  - ***Auto***. VMware picks the fastest and most efficient mode by trying each one on the fly (recommended);
  - ***Advanced***. It's also an auto mode, but here only two high-performance modes are considered to use (Virtual Appliance or SAN). **Virtual Appliance** will be selected if ESX Agent is deployed on a virtual machine hosted by ESX, which guests are being protected or optimized. **SAN** will be selected if ESX Agent is deployed on a physical machine running Windows Server 2003 / 2008 / 2012, which have access to LUNs (Logical Unit Number) with virtual machines. Both modes do not load CPU of ESX. If Virtual Appliance is configured not to use the production networking, then the production bandwidth is also not used. However, both these modes require special configuration of virtual machines.
  - ***Network***. It's a universal mode that can be used no matter where ESX Agent is deployed. It supports all configurations of virtual machines and can process virtual disks that do not support the snapshot mechanism (Independent or Physical RDM disks). Please note however that this mode uses the external network interface of ESX, thus it can heavily load its CPU and bandwidth. Besides, VMware limits for it the number of simultaneously processed disks by all parallel vStrorage API activities (32 disks for ESX 5.x, 27 for ESX 4.x, 23 for ESXi 4.x).
  - ***Encrypted Network***. It's the network mode additionally encrypted through SSL (Secure Sockets Layer). Obviously it's the slowest mode of all. Besides it requires special configuration of virtual

machines. VMware doesn't limit the number of simultaneously processed virtual disks for it. We recommend this mode if ESX Agent connects ESX via public network with low level of confidence.

- ***Simultaneous machine count***. By default, only four virtual machines are allowed to back up simultaneously, which you can change however. Please note the maximum available value is 20 machines.

Advanced settings will be available if the corresponding option is enabled in the Settings dialog.

The second tab (VMs to replicate):

- Browse the connected ESX to mark virtual machines you're going to protect. If guest machines you're going to protect run applications that do not support Microsoft VSS (an old version of MS SQL Server, Linux-based PostgreSQL or Oracle Database, etc.), you need to run custom scripts to provide a coherent state of all open files and databases involved in a backup. Click here for more information.

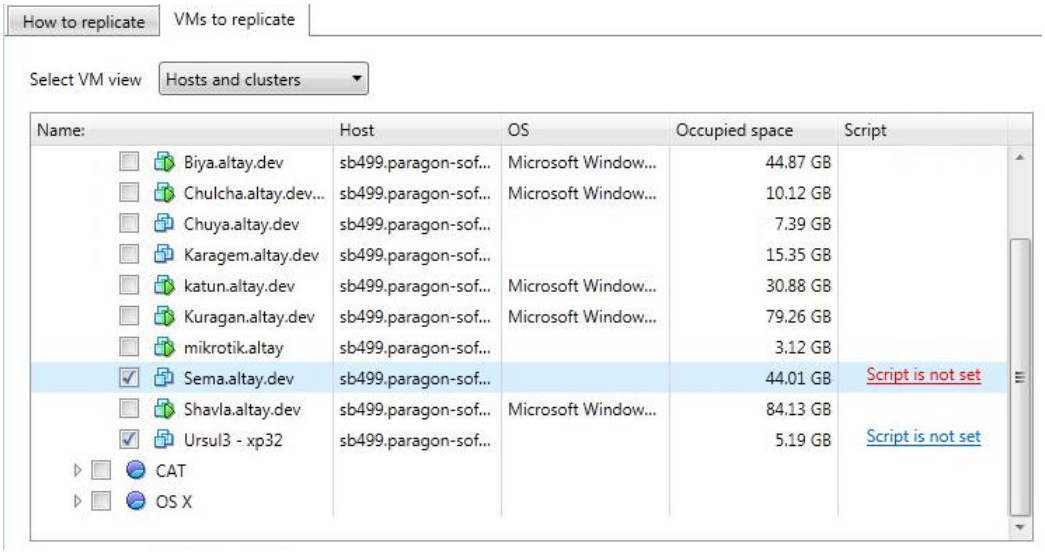

When you're ready with all parameters, click **Save & Validate** to complete creation of the replication policy. By default there will be used the fast level of validation, which you can change by clicking on the arrow button.

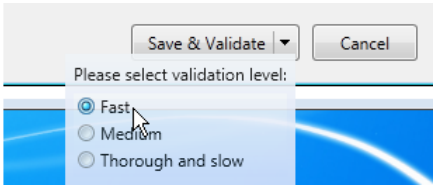

Let's see how three validation levels differ:

- **Fast**. It includes checkup of all policy rules and their parameters, availability of the ESX storage and ESX connection parameters.
- **Medium**. It includes connection to the specified ESX host to scan for target virtual machines as well as connection to the ESX storage to retrieve metadata from it.
- **Thorough and slow**. It includes creation/deletion of snapshots of target virtual machines, creation of an uncompleted replication session and data items in the ESX storage without opening data streams and data copying.

Pre- and post-scripts are also processed during validation launches.

4. Validation of the replica task will be initiated immediately. You will be informed on the operation start through a popup window.

5. Additionally we recommend you to create a notification sub-policy for easier monitoring.
6. If the just created policy hasn't been scheduled, you need to manually commit it once the validation is over. To do that, please go to **Policies > Replication policies for virtual machines**, then select **Run now** for the corresponding policy.

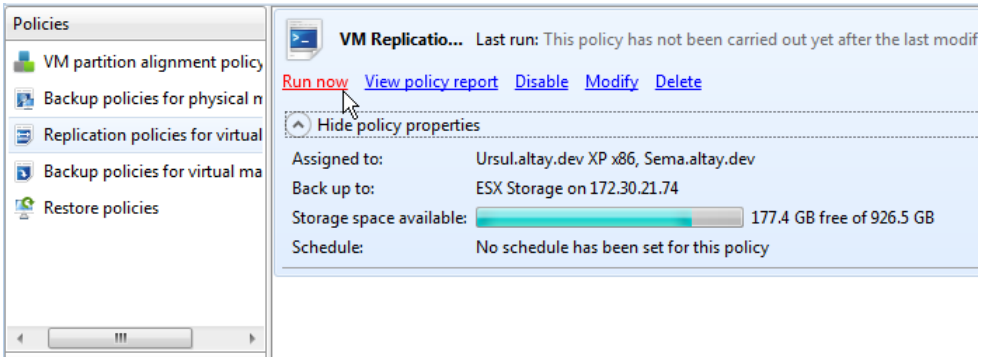

If it's not the first time you commit a backup policy, you will be offered to choose the required replication mode.

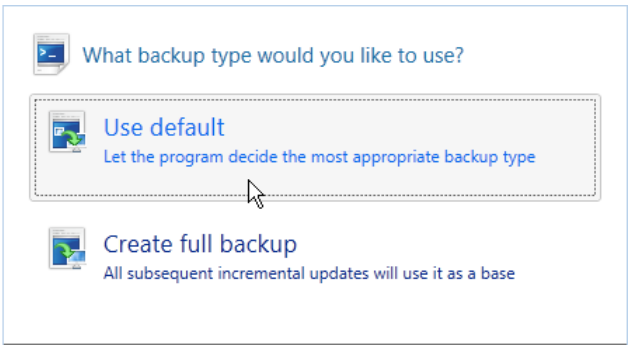

7. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
8. When the replication task is over, its status will be updated.

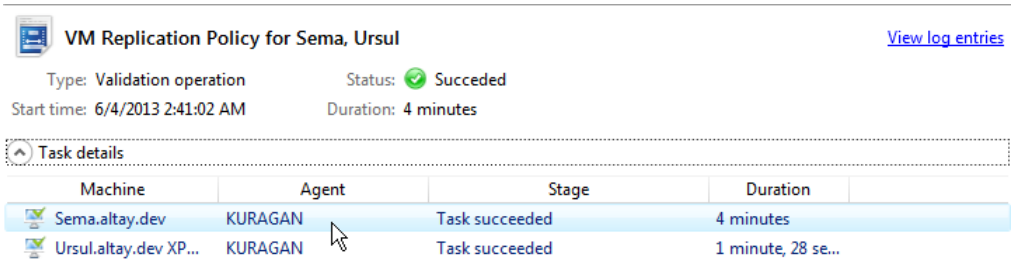

9. If going to the ESX host that accommodates ESX Storage specified for this replication task, you can see two new replicas corresponding to the target virtual machines.

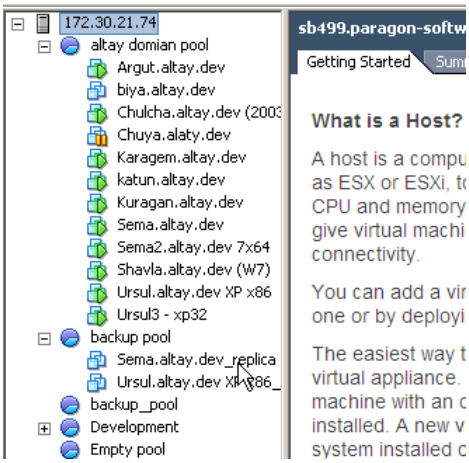

Please do not delete replica virtual machines using the vSphere interface, but only through our consoles. Otherwise you won't be able to do replicas again to the same ESX storage.

To know how to manage created policies, please consult the Managing Policies chapter.

Restoring a VM Backup to a New Location

With PPR you can recover a virtual machine to any good-to-know point in time and place it to the original or a new location. When restored to a new location you will be prompted to provide a new name for the machine, and a host and datastore to reside it. Our product will change the VM configuration file and store the target machine according to the defined location.

To a restore a backed up virtual machine to a new location, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a VM Backup**,

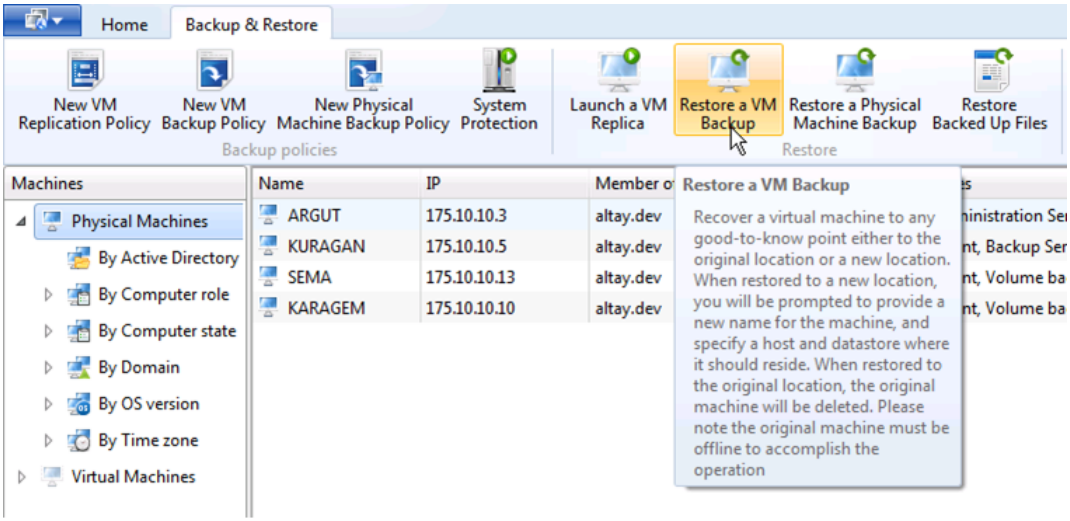

or go to **Policies >** right click on the **Backup policies for virtual machines**, then select **Restore a VM Backup**.

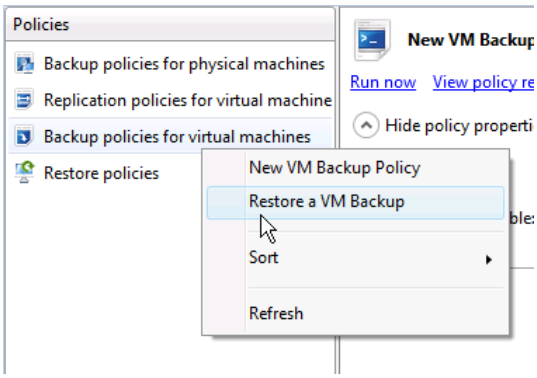

3. The opened wizard will first prompt you to select one of the backed up earlier virtual machines. If there are too many items on the list, please use the search pane to find the required machine by name.

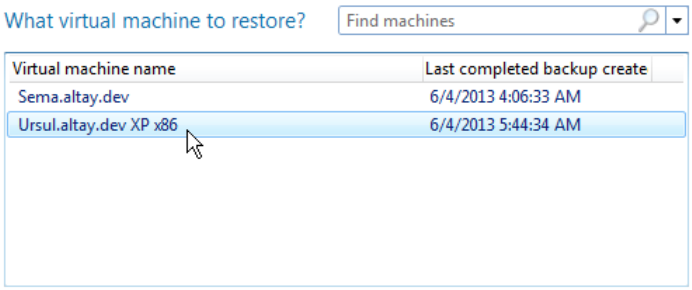

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time. By default, there will be displayed backup images stored in primary backup storages only. If you'd like to see all available backups, please unmark the corresponding option.

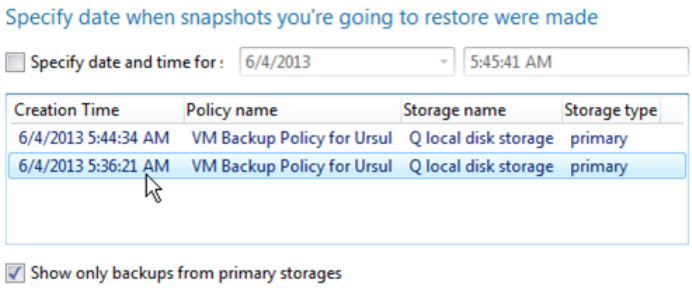

Despite the fact that you're allowed to initiate complete restore or retrieval of certain files/folders from invalid backup images, please do it at your own risk. Please consult the Administering storage backup data chapter to learn how to check images for integrity.

To know how to copy backup data to secondary storages, please consult the Setting up a dual backup strategy chapter.

5. Click on **Restore as a new virtual machine**.

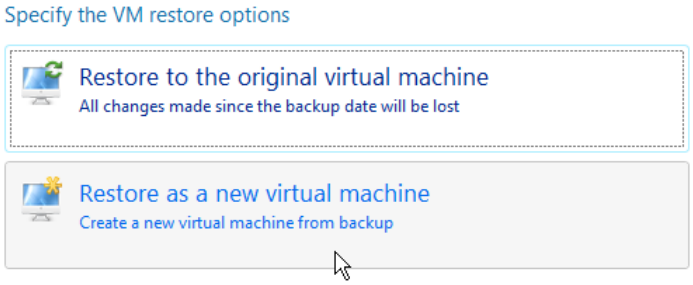

6. Specify an ESX host and credentials where you'd like to deploy this backed up machine. If necessary, set a communication port.

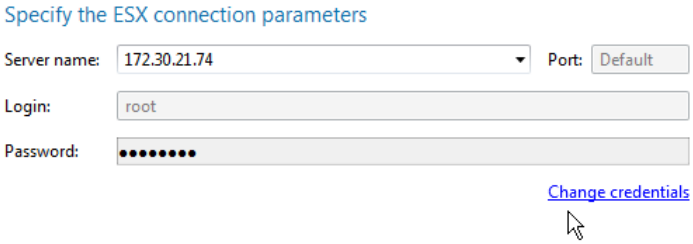

7. Select a resource pool and a datastore, if several.

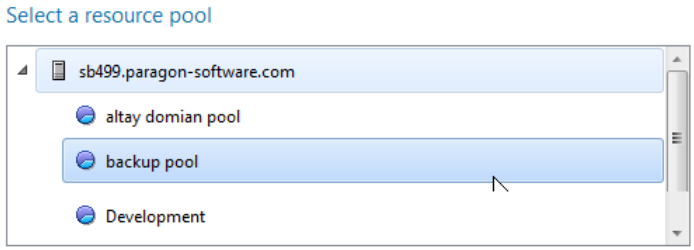

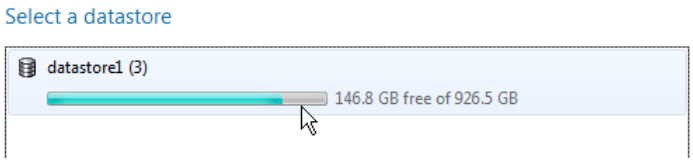

8. Give a name to the new virtual machine and the required provisioning type for its virtual disks. Click **Restore** to initiate the operation.

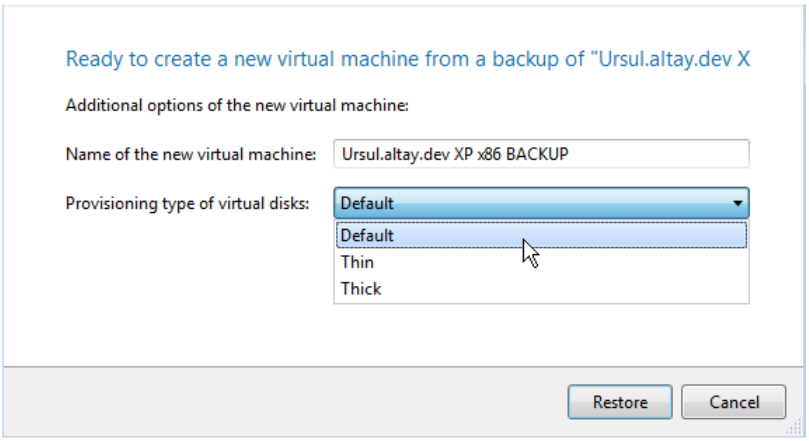

9. You will be informed on the operation start through a popup window.

10. Additionally we recommend you to create a notification sub-policy for easier monitoring.
11. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
12. When the restore task is over, its status will be updated.

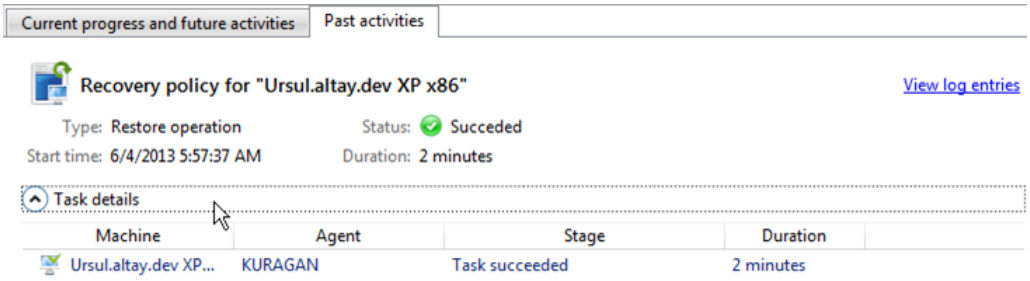

13. If going to the ESX host where you chose to deploy the backed up machine, you can see a new virtual machine in the offline state. You can turn it on, if the original machine is off. Otherwise, there will be a conflict of DNS names or IP addresses (if static addresses are used).

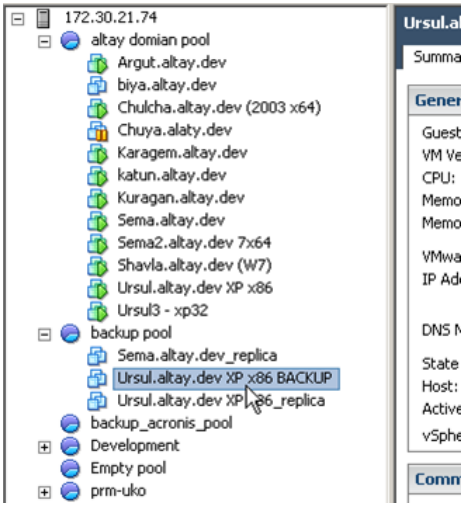

Restoring a VM Backup to the Original Location

With PPR you can recover a virtual machine to any good-to-know point in time and place it to the original location. When restored to the original location, the original machine will be deleted (it should be offline).

To a restore a backed up virtual machine to the original location, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a VM Backup**,

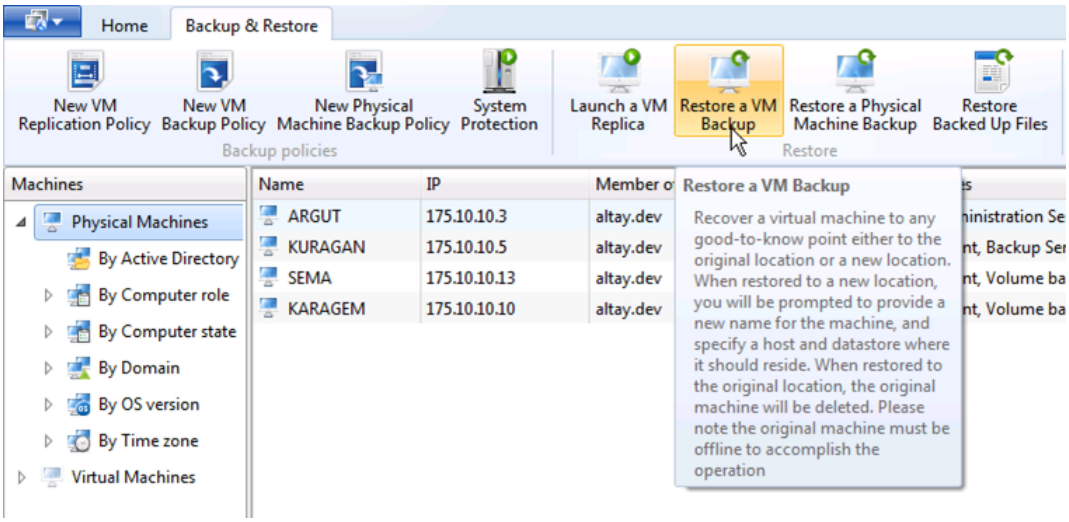

or go to **Policies >** right click on the **Backup policies for virtual machines**, then select **Restore a VM Backup**.

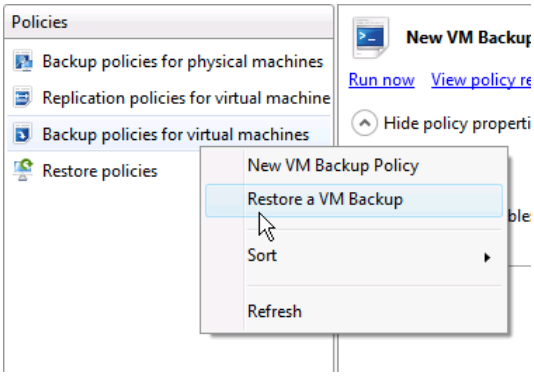

3. The opened wizard will first prompt you to select one of the backed up earlier virtual machines. If there are too many items on the list, please use the search pane to find the required machine by name.

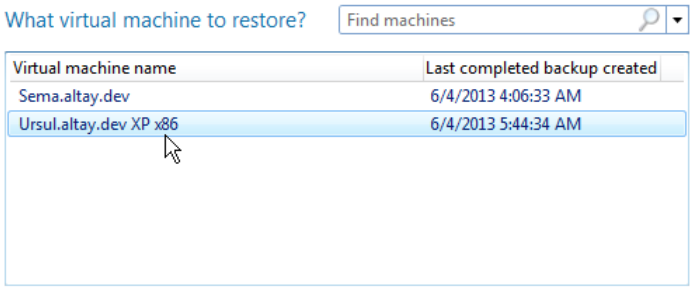

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time. By default, there will be displayed backup images stored on primary backup storages only. If you'd like to see all available backups, please unmark the corresponding option.

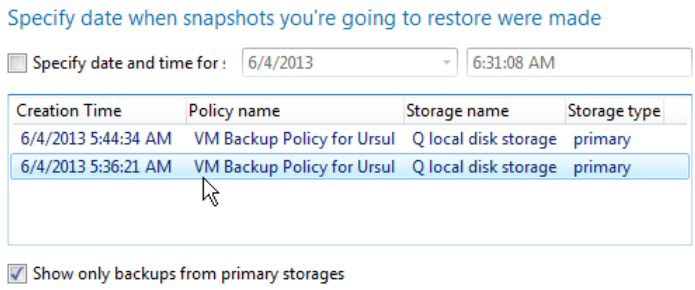

5. Click on **Restore to the original virtual machine**.

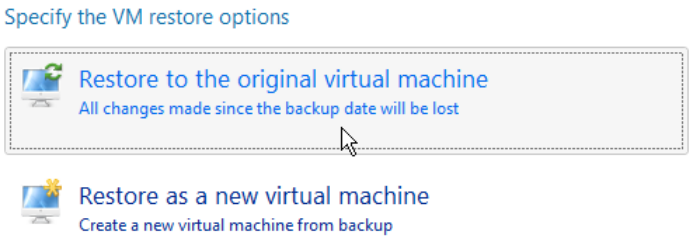

6. You will be informed on the upcoming operation. Please check all parameters are ok. Click **Restore** to initiate the operation.

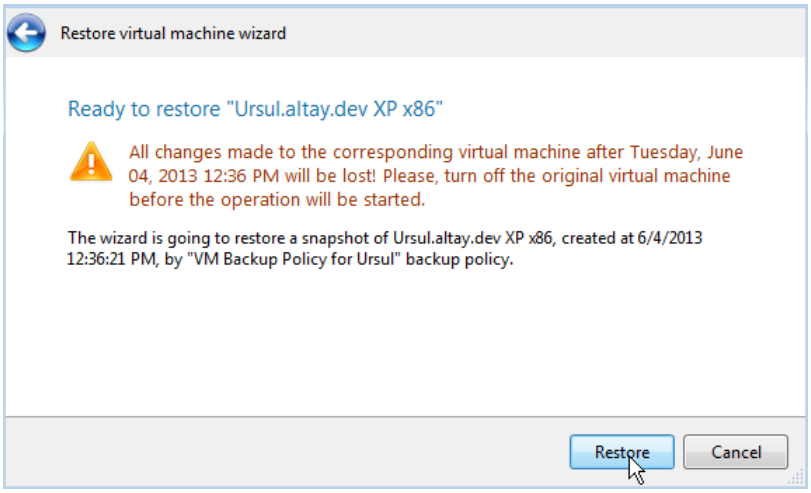

Restore to the original location can only be a success if the target virtual machine is offline.

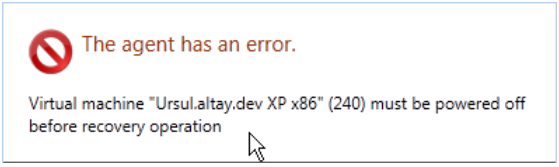

All changes appeared on the target virtual machine after the specified restore point will be irreversibly lost.

---

7. You will be informed on the operation start through a popup window.

8. Additionally we recommend you to create a notification sub-policy for easier monitoring.
9. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
10. When the restore task is over, its status will be updated. Now we can power on the restored virtual machine.

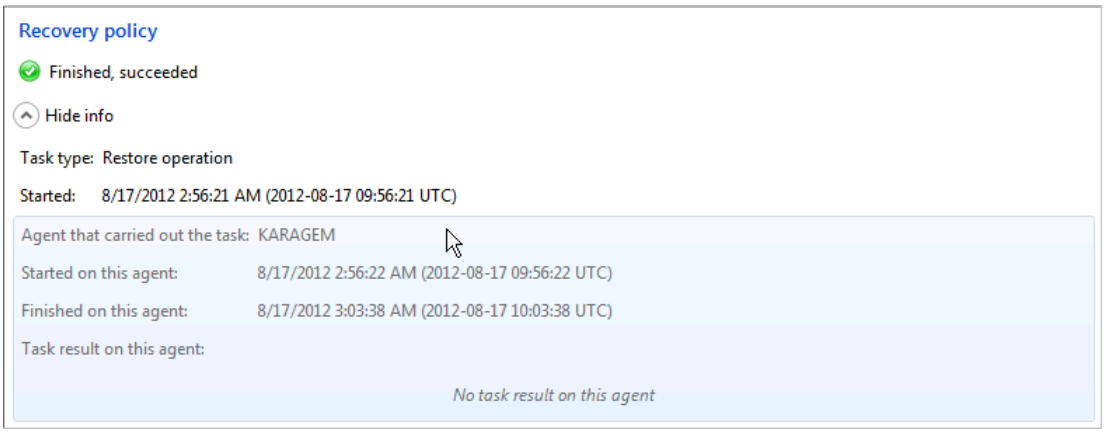

Replica Failover

In case of emergency you can get a problem virtual machine back on track by failing over to one of its replicas, this way a replicated machine takes over the role of the original production machine. You've got the option to fail over to any available time stamp. The whole operation may take only a couple of seconds.

Please do not launch replicas directly from ESX, but use our Replica Failover Wizard. Otherwise, all incremental updates created by our product since the start of the launched this way replica will be corrupted.

Please do not delete replica virtual machines using the vSphere interface, but only through our consoles. Otherwise you won't be able to do replicas again to the specified ESX storage.

To fail over a virtual machine to one of its replicas, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Failover & Launch** ribbon, then select **Replica Failover**.

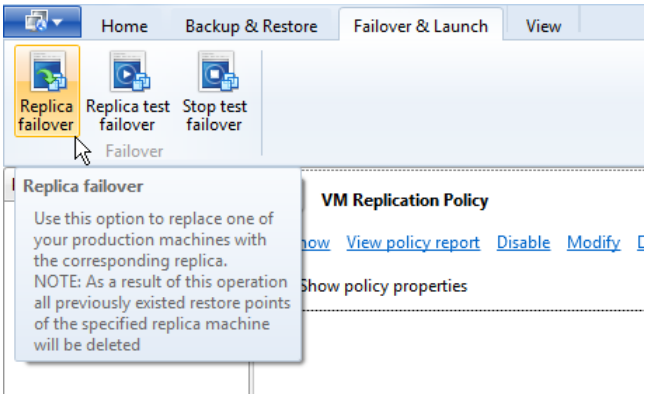

3. The opened wizard will first prompt you to select one of the replicated earlier virtual machines. If there are too many items on the list, please use the search pane to find the required machine by name.

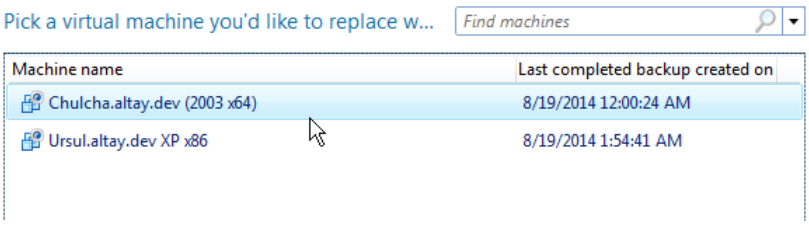

4. Then you will need to choose a desired time stamp, if several. If there are too many items on the list, filter the list by marking the checkbox **Find snapshots created closest to…**, then providing the required date and time.

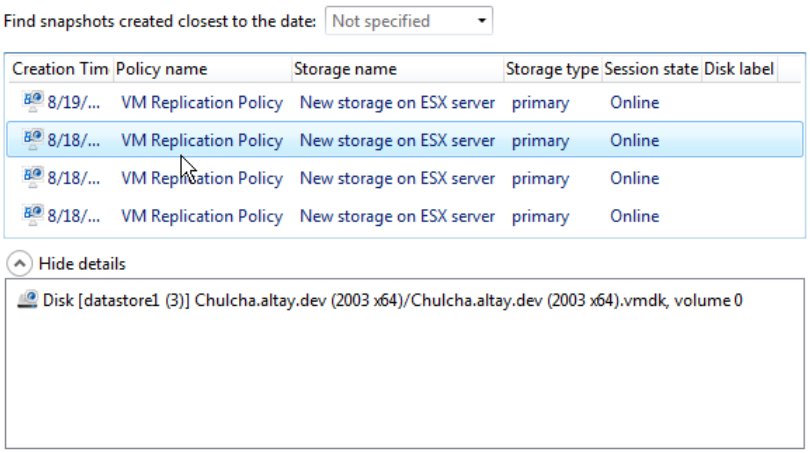

5. You will be informed on the upcoming operation. Click **Finish** to initiate the operation.

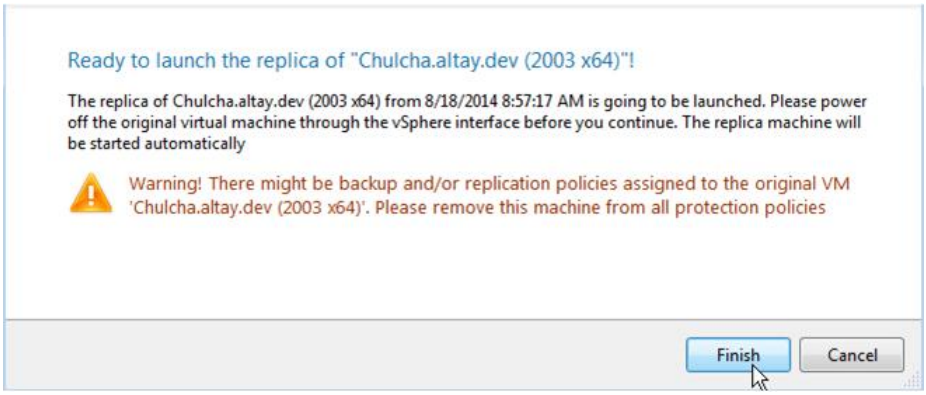

In the current version of the product, the original machine won't be automatically powered off. You should do it manually. Anyway the selected replica will be correctly detached from the replication process and then launched.

Before you continue, please make sure the target virtual machine is not maintained by any backup or replication policy. If it is, please remove it from all backup or replication policies.

Replica Test Failover

This option can help you test the sanity of any time stamp of an existing replica machine, in other words to non-disruptively simulate recovery procedure in an isolated network environment. This operation can be of great use if:

- You'd like to make sure a certain replication policy produces valid replica machines;
- You'd like to do field test for your recovery plan to rely on it in case of disaster;
- You'd like to train your personnel on what is to be done in case of emergency.

Replica Test Failover tasks may take plenty of system resources (CPU, RAM, disk IO), thus please do not forget to stop test replicas when you don't need them.

To test a time stamp of an existing replica machine, please do the following:

1. Launch Protect & Restore Console.

2. If a connection with the server has been established, click on the **Failover & Launch** ribbon, then select **Replica Test Failover**.

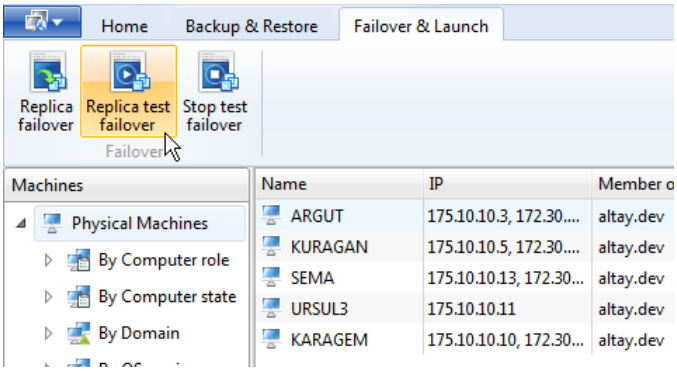

3. The opened wizard will first prompt you to select one of the replicated earlier virtual machines. If there are too many items on the list, please use the search pane to find the required machine by name.

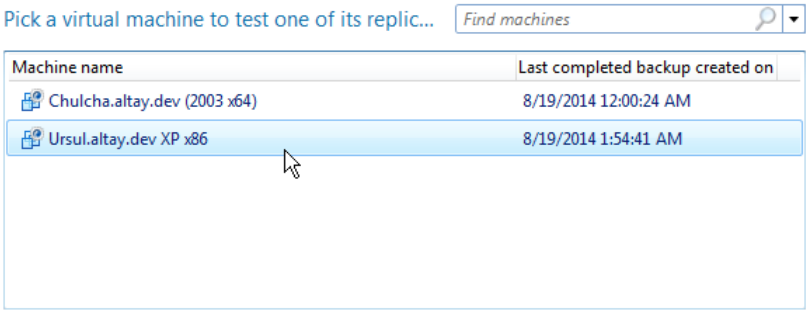

4. Then you will need to choose a desired time stamp, if several. If there are too many items on the list, filter the list by marking the checkbox **Find snapshots created closest to…**, then providing the required date and time.

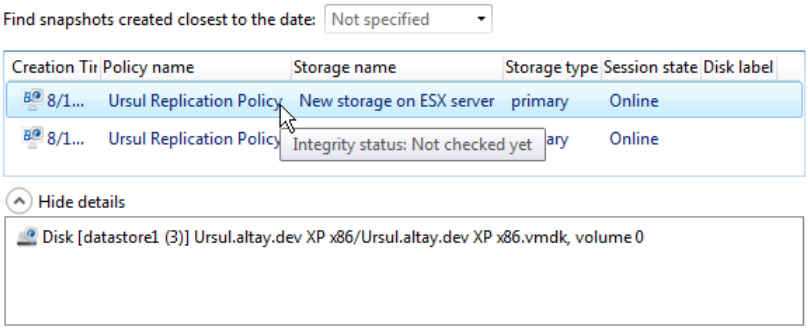

5. By default, the network support on the target replica machine will be disabled to avoid possible problems of having two identical machines in one network environment. If you're using an isolated network, mark the corresponding option to enable the network support. Click **Finish** to initiate the operation.

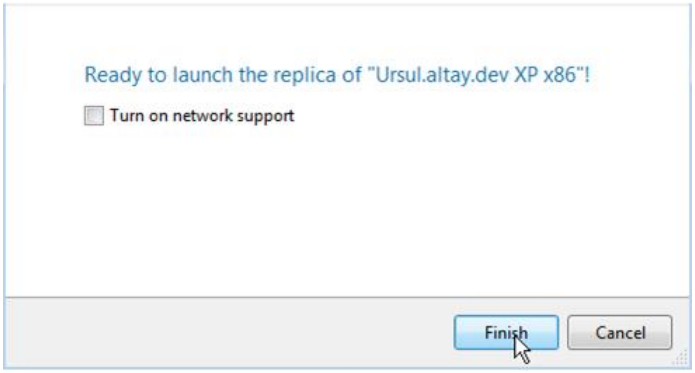

6. When the operation is over, its status will be updated.

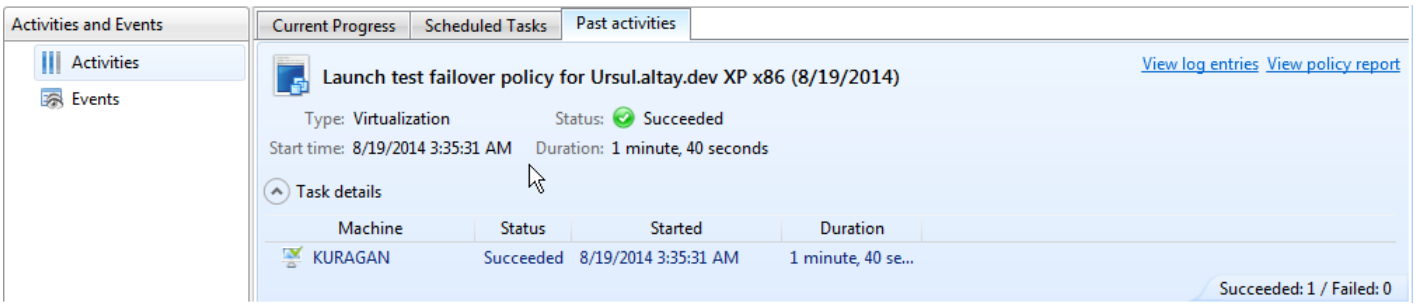

7. You can find the resulted test replica machine in the online state located next to the original replica.

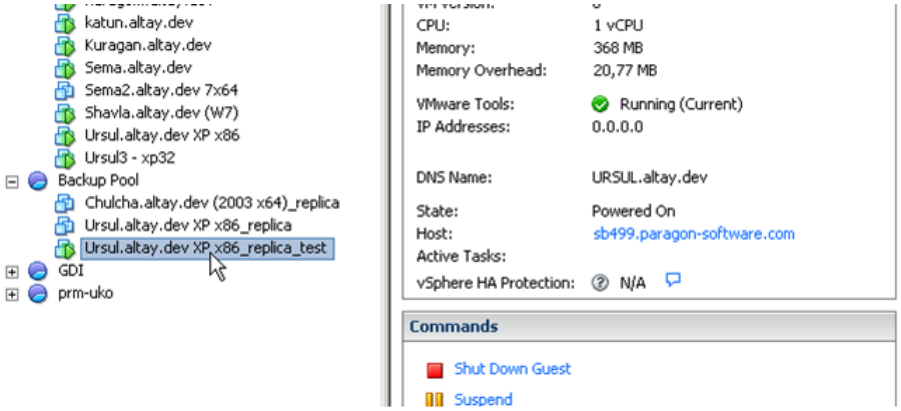

Stop Test Failover

To stop a test failover operation, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Failover & Launch** ribbon, then select **Stop Test Failover**.

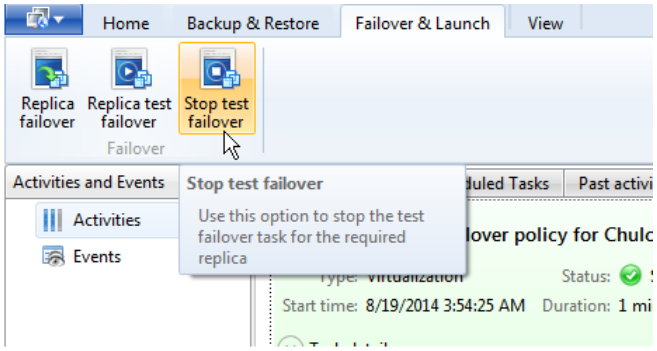

3. The opened wizard will list all replica machines that are currently going through the test failover procedure. Select the required machine, and then click **Stop** to initiate the operation. As a result the selected test replica machine will be turned off and deleted from ESX datastore.

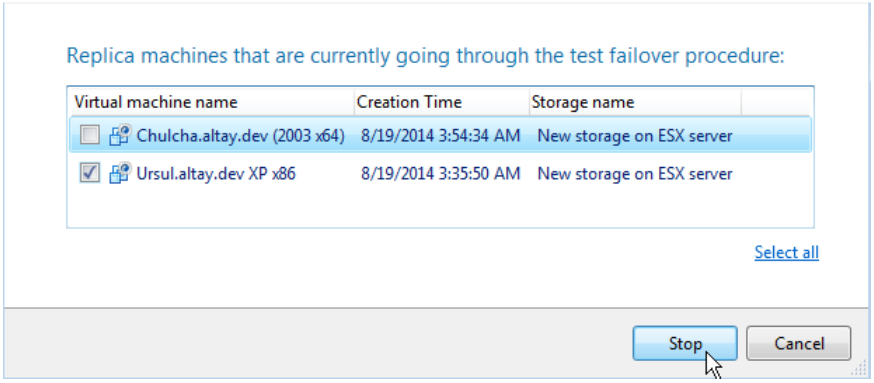

Launching Backup (Instant Restore)

The launch backup aka instant restore is another feature that helps you minimize downtime of a failed production system. It enables to immediately run a physical or virtual machine directly from one of available restore points in VMware ESX environment. Thus users may continue their activities, while you've got enough time to pinpoint and fix the failed system.

When a launch backup operation is performed, PPR initiates creation of an NFS (Network File System) datastore on the specified ESX host, maps the selected backup image to it and then configures a virtual machine that uses the created datastore as disk storage. Since there's no need to extract and copy the image contents to specific location, the whole operation takes a couple of minutes.

The original backup image is locked for writing in order not to break an incremental chain it belongs to, all changes are stored to the NFS datastore. Thus these changes are discarded once the launch backup operation is stopped. It's ok if using this feature for testing purposes just to make sure the target OS and applications are functioning properly. But if you use it in a real disaster recovery scenario you obviously need to save the changes and complete the restore job. You've got several options:

- Migrate the launched machine to production storage through the VMware vMotion technology (no downtime at all);
- Replicate the launched machine to fail over to it when most appropriate (some downtime is inevitable).

At first glance the launch backup has much in common with the replica failover/test failover. Both features are primarily used for high-availability environments that run the first tier applications. However, unlike replication that only works with the VMware- native containers, you can launch any Windows-based physical or virtual machine of any hypervisor out of a backup image, thus opening an easy way for P2V or V2V migration. If comparing operation performance, the launch backup obviously provides limited I/O throughput.

Machines launched out of backup may take plenty of system resources (CPU, RAM, disk IO), thus please do not forget to stop these machines when you don’t need them.

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Backup Server** is installed on any machine, but the more powerful, the better.
- **Backup Virtualizer plug-in** is added to Backup Server.
- There should be configured a backup storage containing at least one backup image.
- ESX host where you’re going to launch backup images should resolve Backup Server by name.

Operation scenario

To launch a physical or virtual machine out of a backup image, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Failover & Launch** ribbon, then select **Launch Backup**.

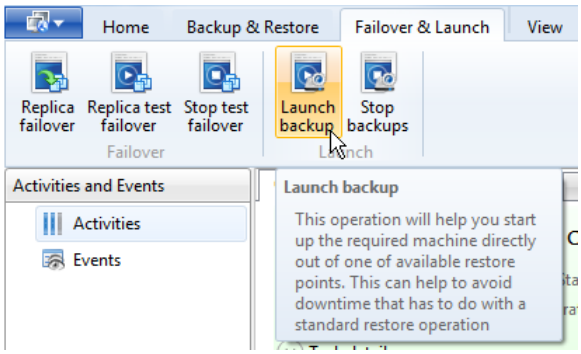

or go to **Infrastructure > Backup Servers**, select the required storage, then click the **Browse storage** icon to see machines it contains. Select the required restore point, then click **Launch** to initiate the operation.

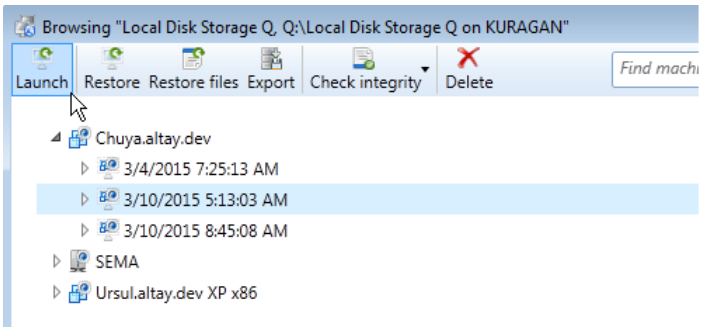

If your Backup Server doesn't have the Backup Virtualizer plug-in, you will be prompted to install it.

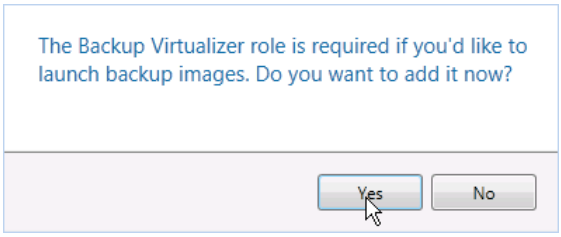

3. The opened wizard will first prompt you to select one of the previously backed up machines. If there are too many items on the list, please use the search pane to find the required machine by name.

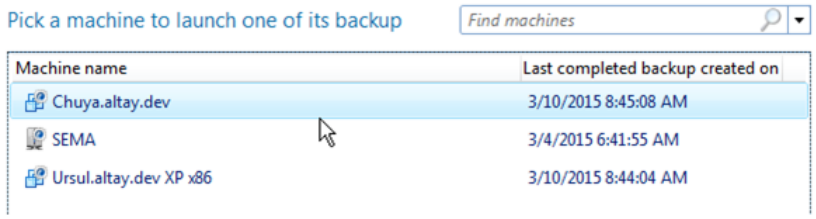

Beside ESX guest machines, PPR enables to launch any Windows-based physical or virtual machine out of a backup image.

4. Then you will need to choose a desired time stamp, if several. If there are too many items on the list, filter the list by marking the checkbox **Find snapshots created closest to…**, then providing the required date and time.

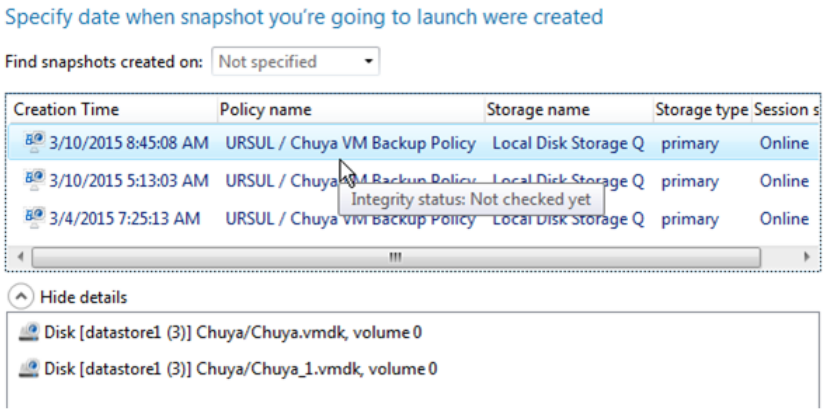

5. Enter a DNS name or IP address of the required vCenter or ESX host, a communication port (if necessary), and administrator credentials in the corresponding fields.

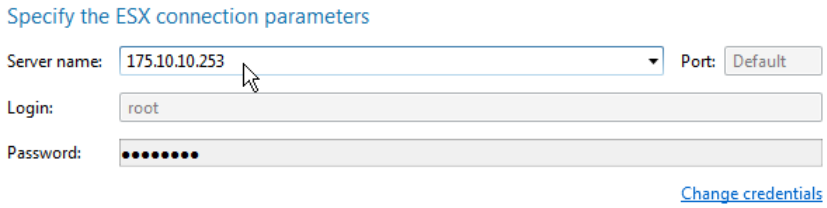

If the required ESX host is a member of a vCenter, always use the IP address and credentials of that vCenter.

6. If the provided IP and access credentials are valid, there will be established connection to the specified VMware infrastructure. Select a resource pool to place the launched machine to.

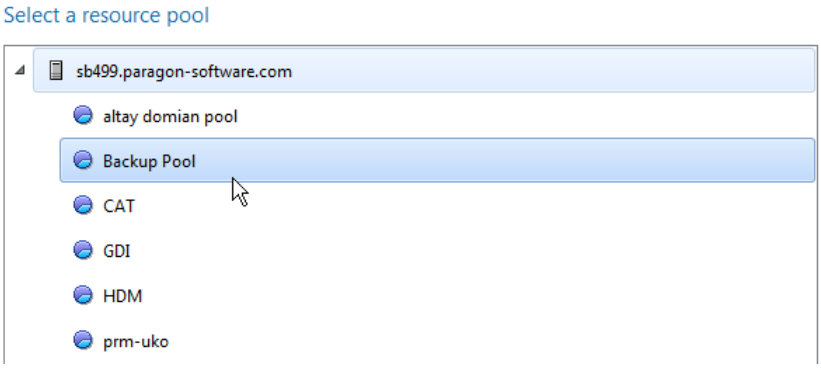

7. By default, the network support on the target machine will be disabled to avoid possible problems of having two identical machines in one network environment. If you're using an isolated network, mark the corresponding option to enable the network support. At this stage you can also change the offered machine name. Click **Finish** to initiate the operation.

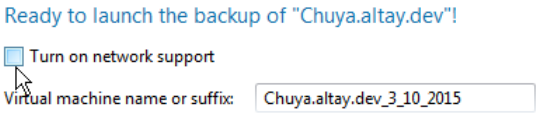

8. When the operation is over, its status will be updated.

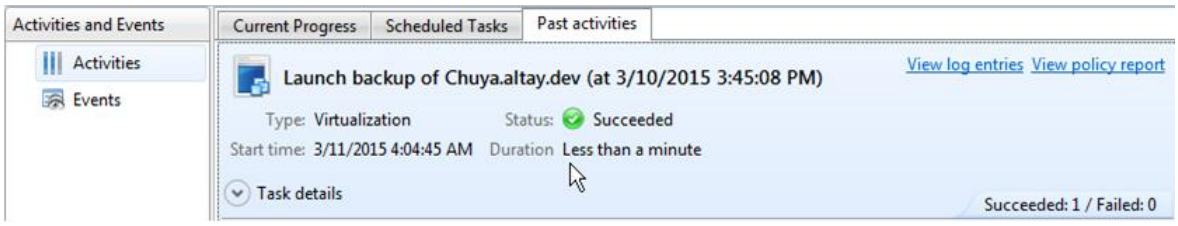

9. You can find the resulted machine in the online state in the specified resource pool.

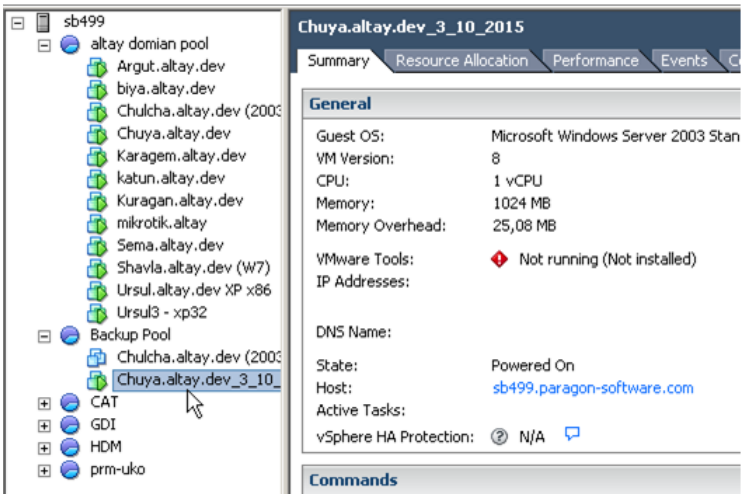

Stop Launch Backup

To stop a launch backup operation, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Failover & Launch** ribbon, then select **Stop Backups**.

3. The opened wizard will list all currently launched machines. Select the required machine, and then click **Stop** to initiate the operation. As a result the selected machine will be turned off and removed.

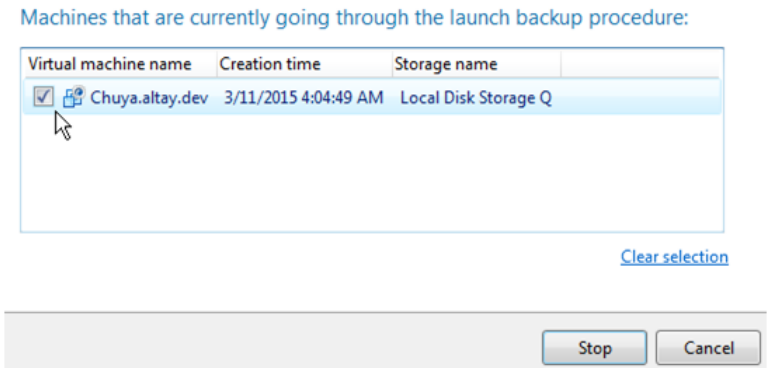

Restoring Separate Files

PPR allows browsing contents of virtual or physical backup images as well as VM replicas to do granular recovery of separate files and/or folders. Required data can be restored either locally (on a machine where Protect & Restore Console is installed) or on a network share, provided the original directory structure is kept intact if necessary.

To restore separate files and/or folders, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore Backed Up Files** ,

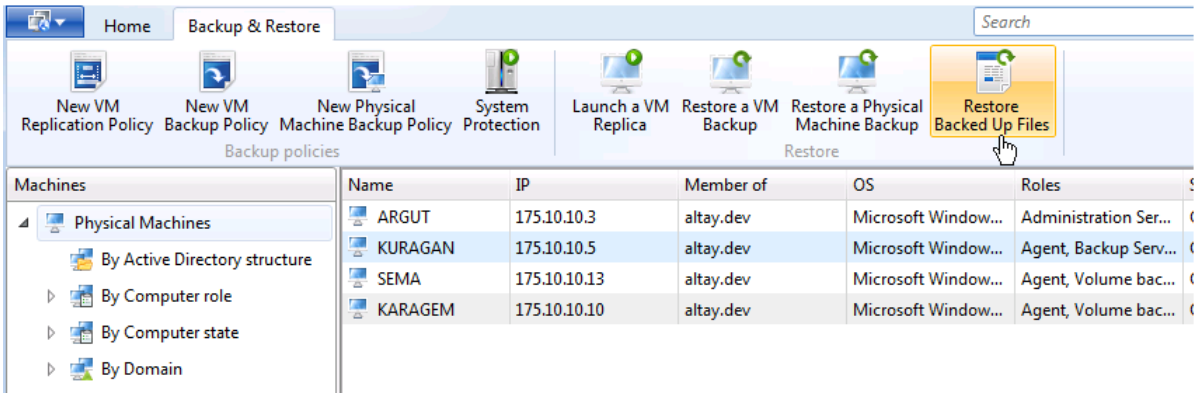

3. The opened wizard will first prompt you to select one of the backed up or replicated earlier machines. If there are too many items on the list, please use the search pane to find the required machine by name.

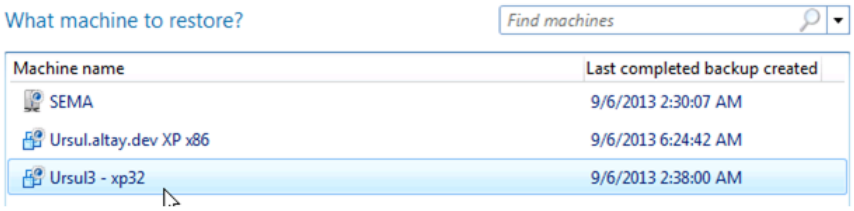

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

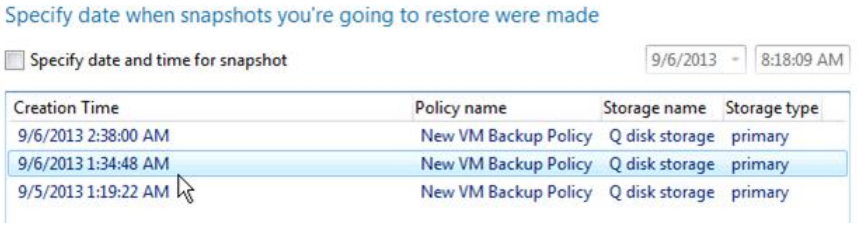

Despite the fact that you're allowed to initiate complete restore or retrieval of certain files/folders from invalid backup images, please do it at your own risk. Please consult the Administering storage backup data chapter to learn how to check images for integrity.

5. Find and mark files and/or folders, you'd like to restore.

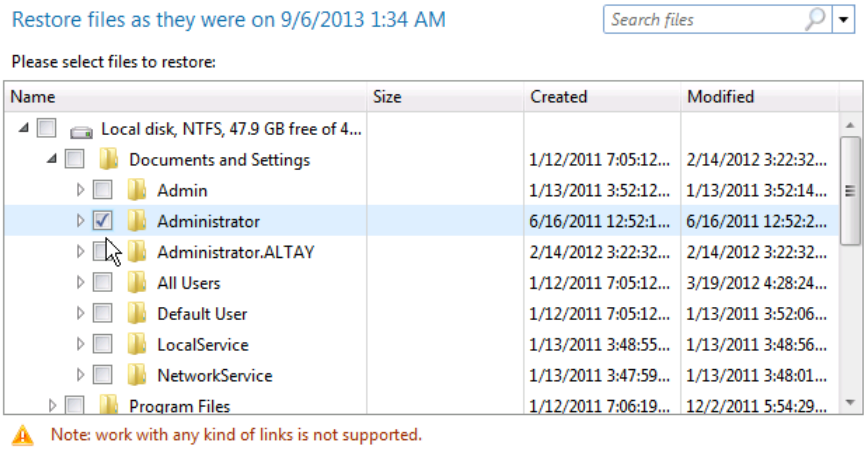

6. Specify where you'd like the selected backup data to be placed to (a local folder of a machine where Protect & Restore console is installed or a network share). If you'd like Windows Explorer to open in the specified folder once the operation is over, please additionally mark the corresponding option. Use the **Preserve directory structure** option to keep the original directory structure intact. Click **Start** when ready.

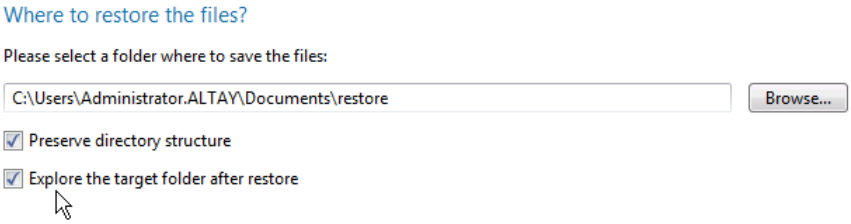

7. Monitor the operation progress. Click **Finish** when it's over.

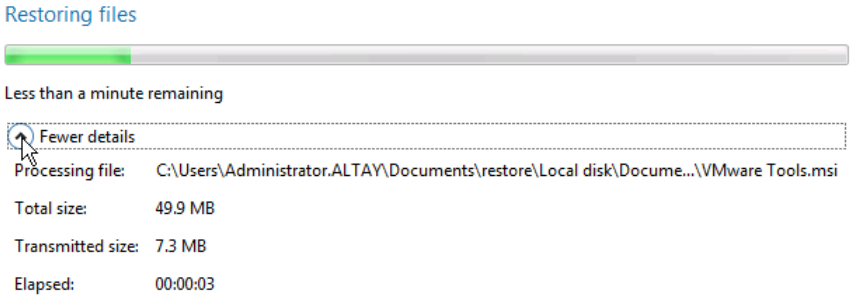

File-level restore is available for all types of backup images and VM replicas.

Protecting Physical Machines

Backing up Physical Machines

PPR allows agent based protection of any physical (or virtual treated as physical) Windows machine (since Windows XP). One backup task can involve one or many machines. When setting up a physical backup policy, you can specify as a backup object entire computers or separate volumes. By default, for every machine our product creates a full backup in a special proprietary format during the first run, then incremental updates according to a set timetable. It allows configuring general retention policies for backup storages or a particular policy for a certain backup task, specifying how long backups should be kept or the amount of space they can take. When time comes, all restore points beyond the set

limit are merged with their full backup thus creating a new full backup. All backup images are being highly compressed during creation by using redundant data exclusion filters (OS page files, zero data blocks, etc.) and a pVHD backup format, which eases the backup storage requirements.

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Backup Server** is installed on any machine, but the more powerful, the better.
- There has been registered a primary local or network backup storage. To know more on the subject, please consult the Registering primary storages chapter.
- **Target machines have been added to the infrastructure**.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **New Physical Machine Backup Policy**,

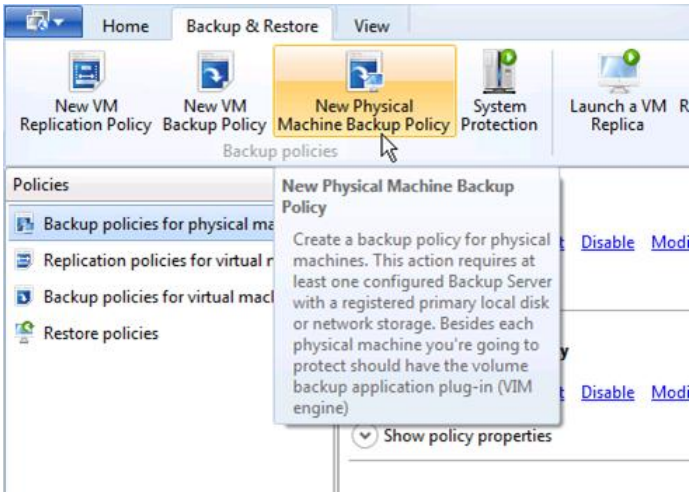

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **New Physical Machine Backup Policy**.

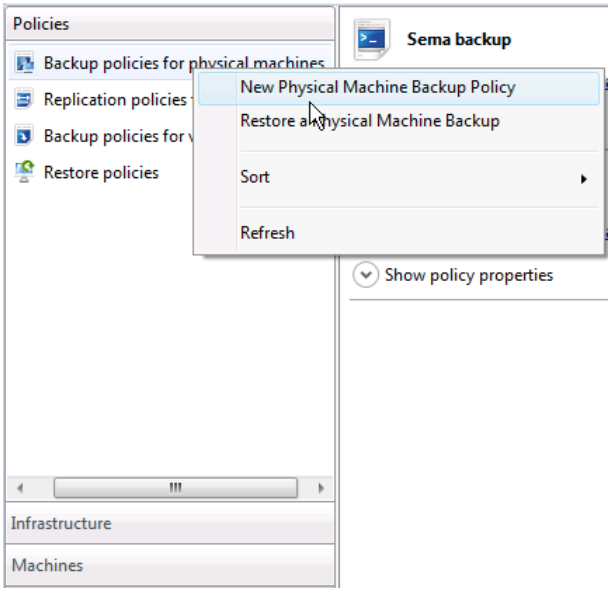

3. The opened dialog consists of four tabs that include a number of parameters:

   The first tab (Policy settings):

   - **Policy name**. Give it a catchy name.

     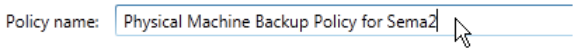

   - **Description**. Give a detailed description to the backup task (optional).

     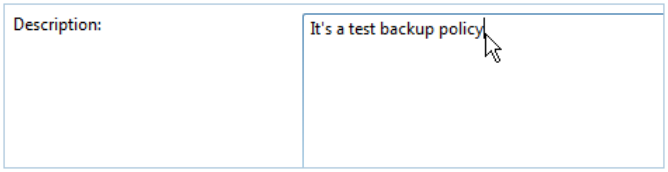

   - **Back up to**. Select a backup server (if several), then the required primary storage from the popup list to place backup images to.

     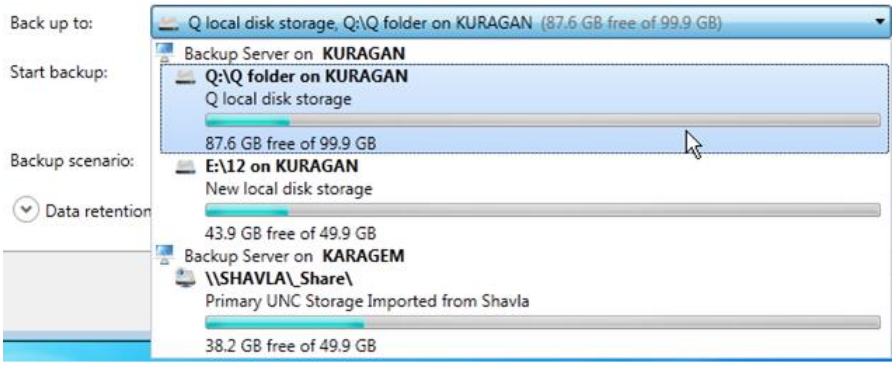

   - **Start backup**. By default, no schedule is set for the backup policy, so you will need to manually commit it after its validation. If you want to schedule the policy, just click on the corresponding link to specify a timetable.

     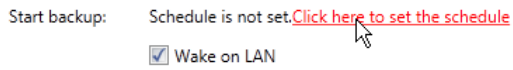

     The opened dialog consists of two sections:

     Basic scheduling

     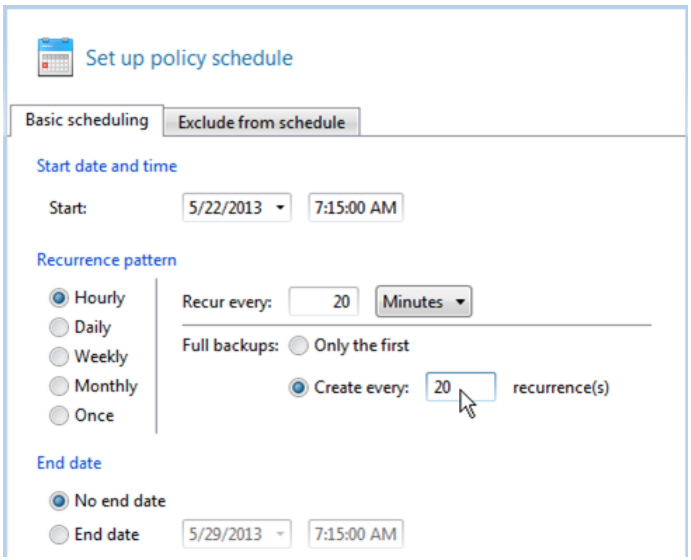

In this section you can set up a backup timetable. By default, a full backup will be created once for every target machine, then only come incremental updates, which you can change however through the **Full backups** section.

Exclude from schedule

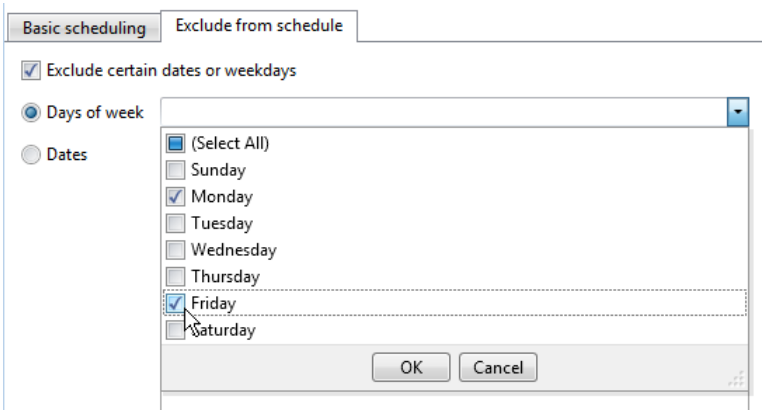

In this section you can specify days of week, or certain dates, when backup operations should not be accomplished.

If you schedule a backup task, the operation will start on each target machine according to its local time.

- **Wake on LAN**. Be default, the target physical machines will be automatically turned on to do backup through the Wake-on-LAN assistant.

- **Backup scenario**. In the current version of the product only one backup scenario is supported (simple).

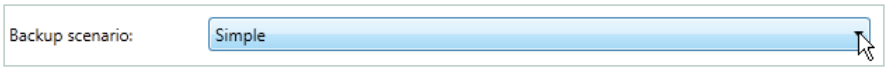

- **Data retention options**. Here you can specify a custom backup data retention mode that will be taken into account for the created policy only.

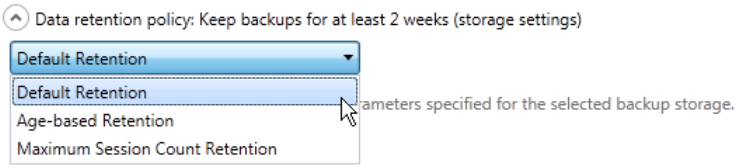

  - **Default Retention** to use data retention parameters specified for the selected backup storage.
  - **Age-based Retention.** Use this option if you'd like to limit lifetime of backup images created by this policy (2 weeks by default, set in the **Age-based criterion** option). To minimize load on the backup server, there's a conditional criterion (**Size-based criterion**) you can make use of to suppress the data retention process until size of backups per each machine exceeds a certain value (10 GB by default).

- **Maximum Session Count Retention** to define the maximum number of backup sessions allowed for target machines processed by this policy. On exceeding the set value, backup chains will be thinned out starting from the oldest backup images.

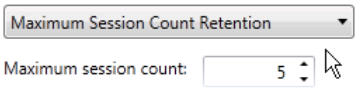

The second tab (Policy objects):

- **Back up entire computer**. Select this option to protect all volumes of all specified target computers.
- **Select objects to back up**. By default, entire machines will be protected. However you can specify particular volumes that need protection. All the rest won't be processed during the backup operation, which can help you minimize the backup storage requirements.

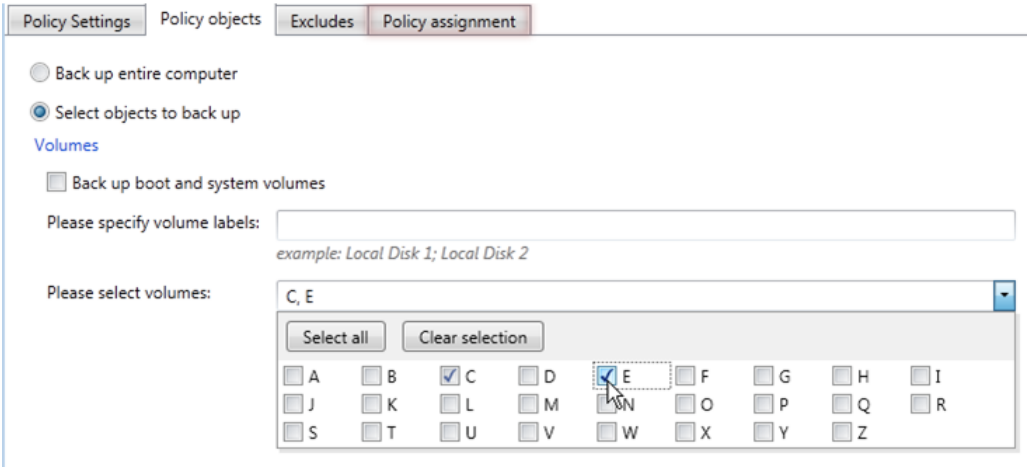

- **Back up boot and system volumes**. Use this option to allow our program to automatically detect system and boot volumes of every target machine and add them to the backup task.
- **Selecting volumes**. You can manually specify volume letters or volume labels that require protection.
- **Back up volumes without drive letters**. Use this option to allow our program to automatically detect and protect volumes that have not acquired drive letters in target operating systems.

Specified in this section parameters will be applied to all target computers.

We do not recommend you to back up different volumes of one and the same machine by different backup policies that all use one and the same backup storage, as in this case only full backup images will be created. However, if several backup policies are configured to protect an identical set of volumes, than the incremental imaging is supported.

The third tab (Excludes). Here you can specify what data should be automatically ignored during backup. You can filter certain files or folders by creating masks. There are two types of filters:

- **Shared** that are applied to all physical backup policies. Click here for more information.
- **Private** that are created and applied to the current backup policy only.

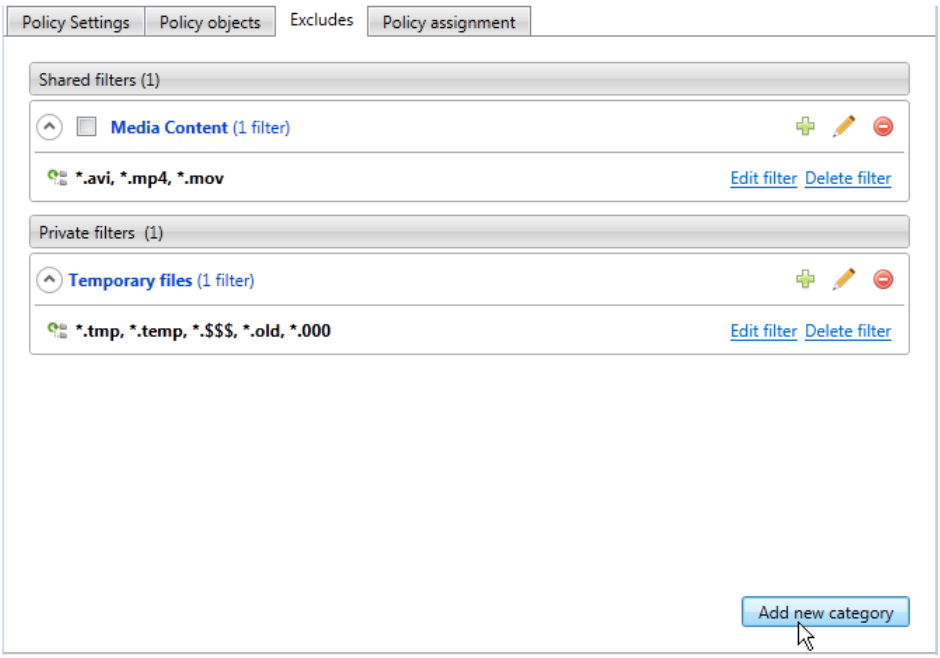

Exclude filters assigned at this stage will be applied to all target machines processed by this policy.

The fourth tab (Policy assignment). In this section you should specify target physical machines you're going to protect. Click on the **Show advanced settings** link next to the specified machine(s) to enable the VSS (Microsoft Volume Shadow Copy Service) logging that can help you pinpoint backup failures. Please note that VSS logs will take plenty of disk space, thus use this option for troubleshooting only.

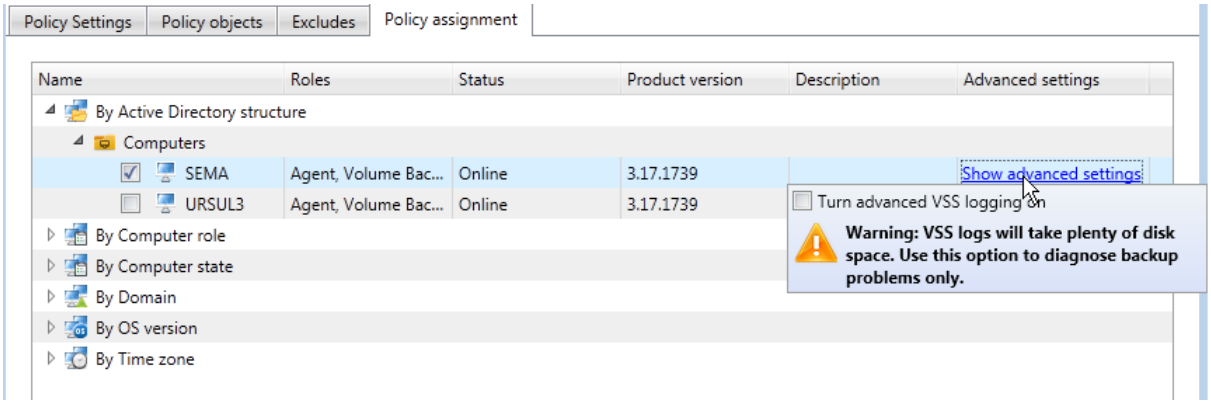

You can only specify machines that are already members of the infrastructure having the Volume backup application plug-in installed.

Advanced settings will be available if the corresponding option is enabled in the Settings dialog.

When you're ready with all parameters, click **Save & Validate** to complete creation of the backup policy. By default there will be used the fast level of validation, which you can change by clicking on the arrow button.

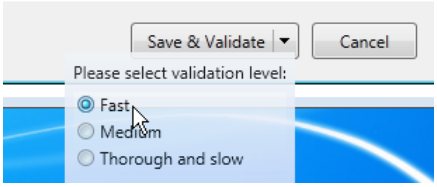

Let's see how three validation levels differ:

- **Fast**. It includes checkup of all policy rules and their parameters, availability of the backup storage, volume backup application and engine plug-ins.
- **Medium**. It includes search of meta items, connection to the specified backup server to get the required backup storage, and retrieval of metadata from the storage.
- **Thorough and slow**. It includes creation/deletion of VSS snapshots of target virtual machines, creation of an uncompleted backup session and data items in the backup storage without opening data streams and data copying.

4. Validation of the backup task will be initiated immediately. You will be informed on the operation start through a popup window.

5. Additionally we recommend you to create a notification sub-policy for easier monitoring.
6. If the just created policy hasn't been scheduled, you need to manually commit it once the validation is over. To do that, please go to **Policies > Backup policies for physical machines**, then select **Run now** for the corresponding policy.

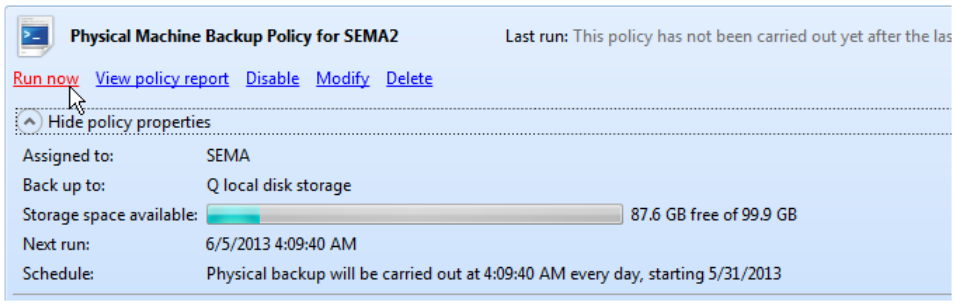

If it's not the first time you commit a backup policy, you will be offered to choose the required backup mode.

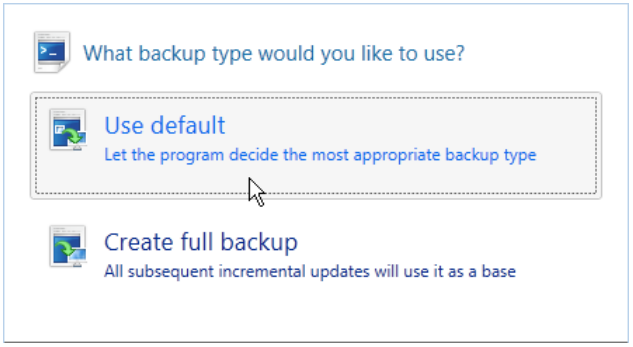

7. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
8. When the backup task is over, its status will be updated.

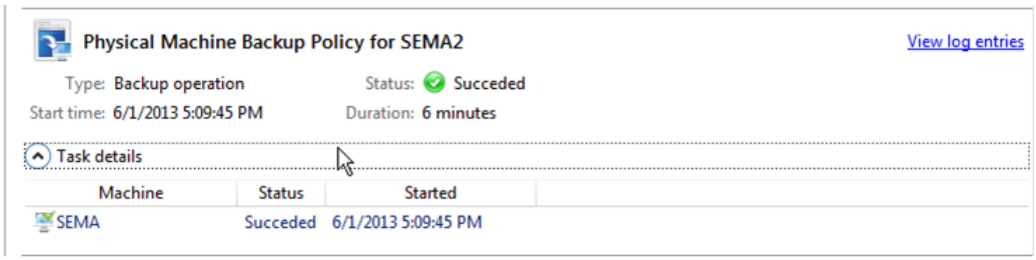

To know how to manage created policies, please consult the Managing Policies chapter.

Immediate Protection of a Stand-alone Physical Machine

PPR includes a special wizard that helps to set up all necessary components of the infrastructure (if not done yet) and carry out other actions required for protection of a single physical machine.

Prerequisites

- **Protect & Restore Server** (Administration and Installation Servers) is installed on any domain machine, but the more powerful, the better.
- **Protect & Restore Console** is installed.
- **Protect & Restore Backup Server** is installed on any machine, but the more powerful, the better.
- **Target machines have been added to the infrastructure**.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **System Protection**.

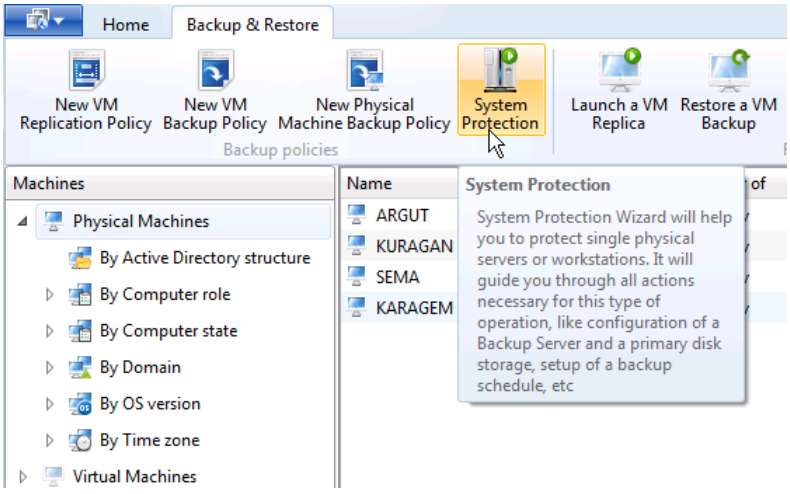

3. The opened wizard will first prompt you to select a machine you'd like to protect.

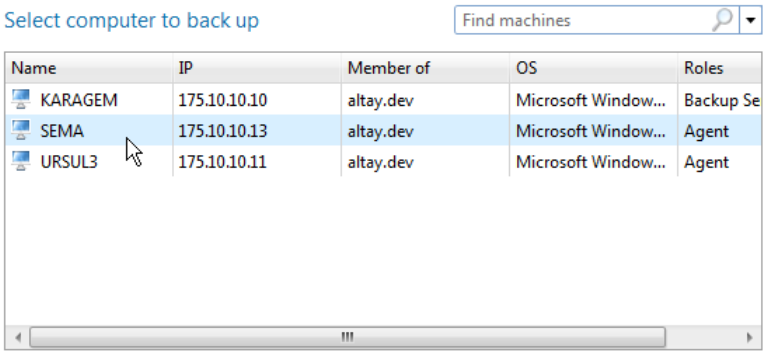

4. Next step you should decide whether to back up the entire machine (all volumes, service information, etc.) to later accomplish bare metal recovery, or only certain volumes.

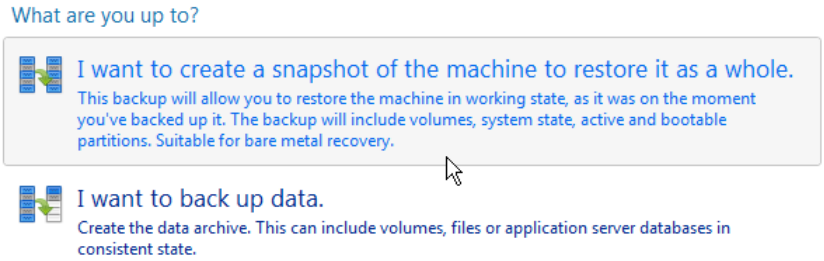

5. On the **Protection & Integrity** tab, you can edit the default backup policy name and specify a custom backup data retention mode if necessary.

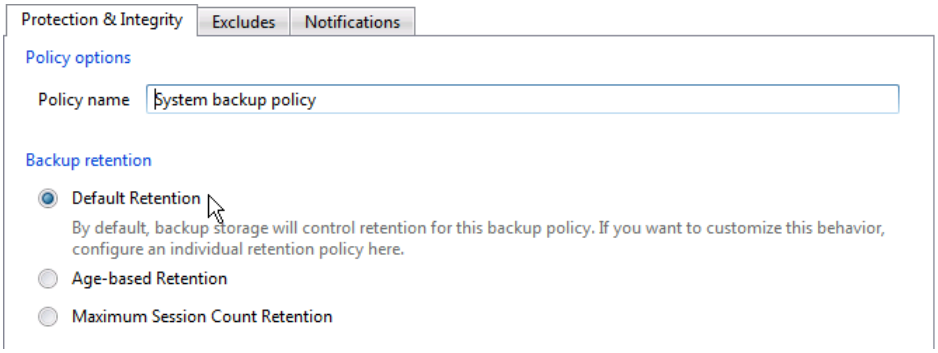

- **Default Retention** to use data retention parameters specified for the selected backup storage.
- **Age-based Retention.** Use this option if you'd like to limit lifetime of backup images created by this policy (2 weeks by default, set in the **Age-based criterion** option). To minimize load on the backup server, there's a conditional criterion (**Size-based criterion**) you can make use of to suppress the data retention process until size of backups per each machine exceeds a certain value (10 GB by default).

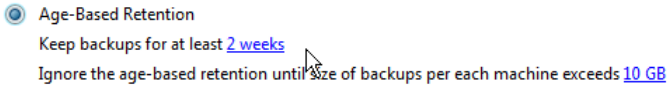

- **Maximum Session Count Retention** to define the maximum number of backup sessions allowed for the target machine. On exceeding the set value, backup chains will be thinned out starting from the oldest backup images.

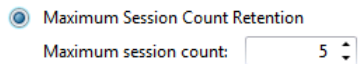

6. Click on the **Excludes** tab to specify what data should be automatically ignored during backup. You can filter certain files or folders by creating masks. Click here for more information.

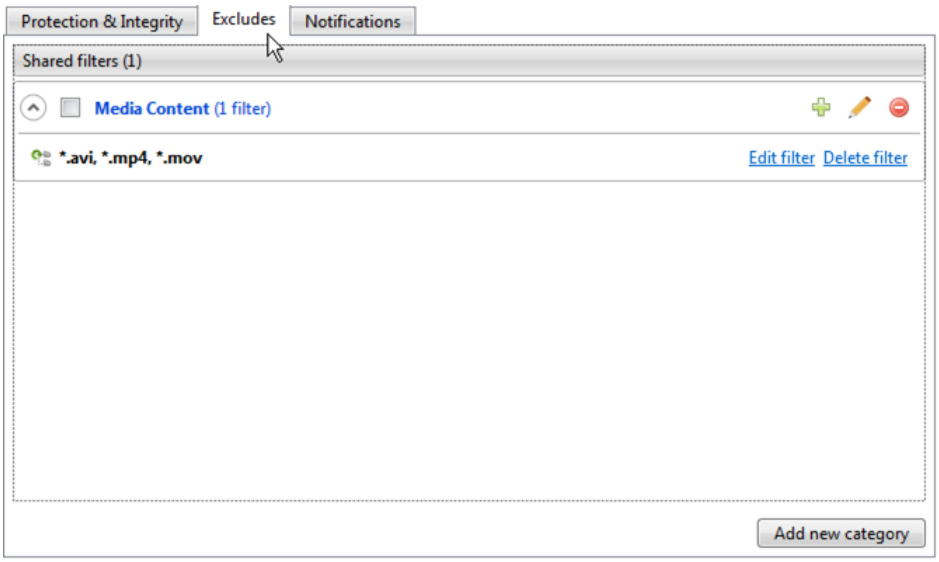

7. Click on the **Notifications** tab to set up a notification policy if necessary.

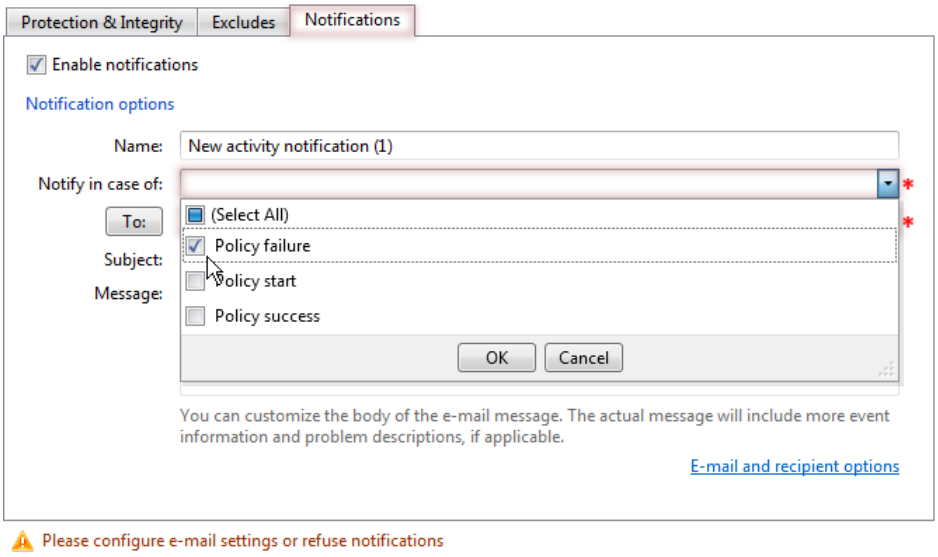

8. Choose the required type of backup storage (local, network, or disk pool), where you'd like the resulted backup images to store.

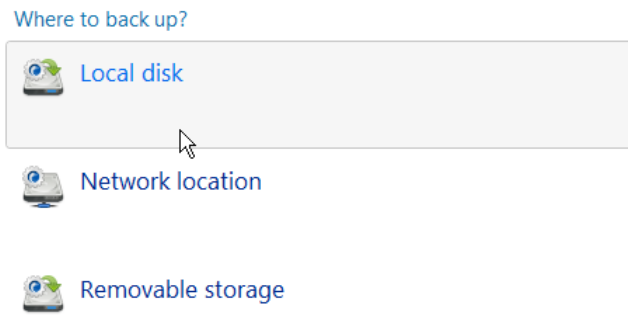

9. Select a backup storage of the specified type from the list of already registered in the infrastructure or register a new primary backup storage on one of the infrastructure backup servers.

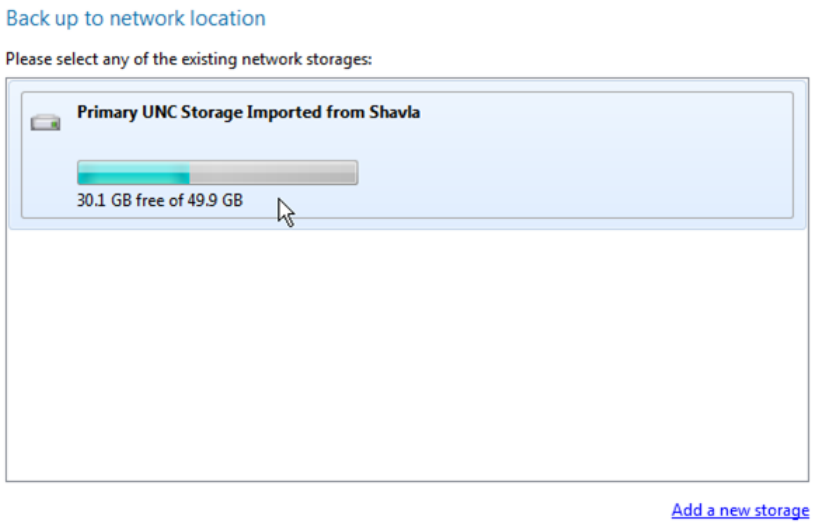

10. Set an operation schedule or launch the backup immediately by choosing one the appropriate options.

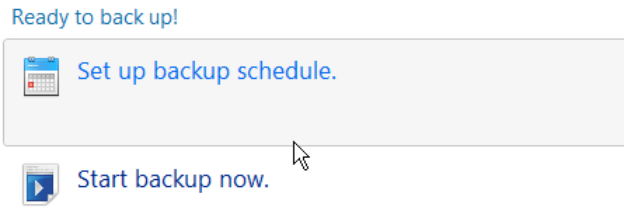

11. Give a catchy name to the resulted backup policy, then check and modify the data retention parameters if necessary. Click **Back up** to initiate the operation.

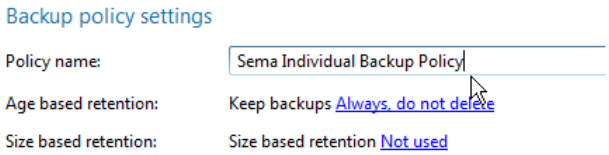

12. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
13. When the backup task is over, its status will be updated.

To know how to manage created policies, please consult the Managing Policies chapter.

Restoring Non-system Volumes Remotely

Restore of data (non-system) volumes can be accomplished remotely through Console.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a Physical Machine Backup**,

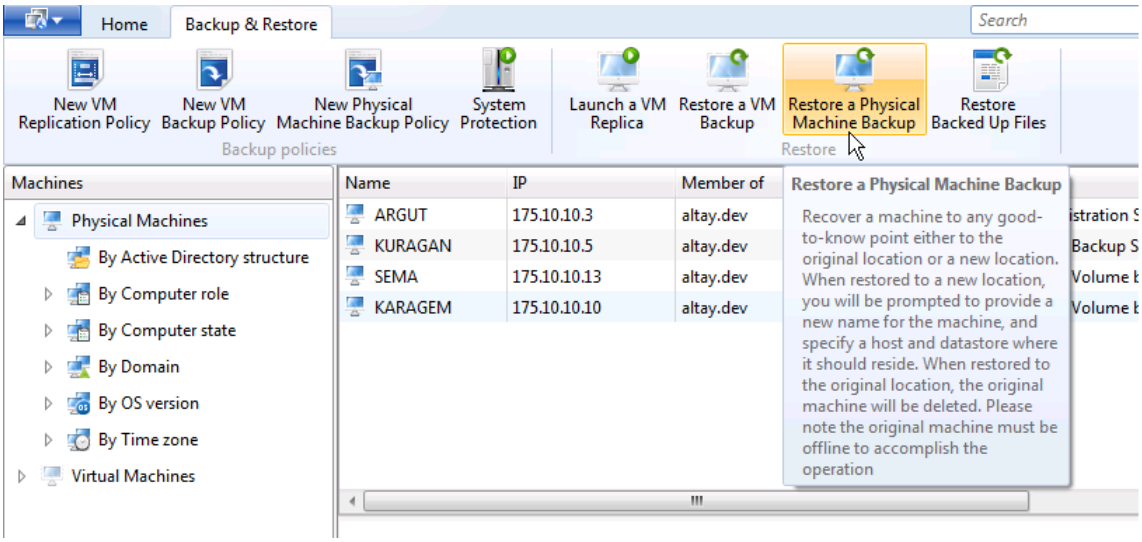

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a Physical Machine Backup**.

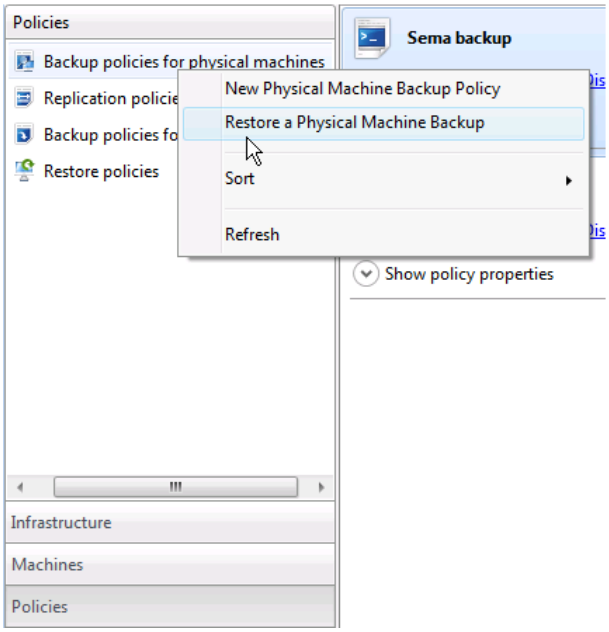

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name.

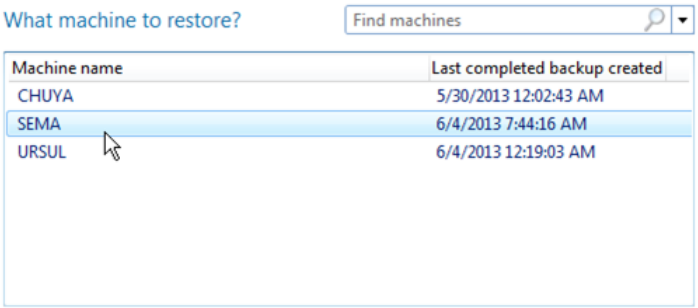

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

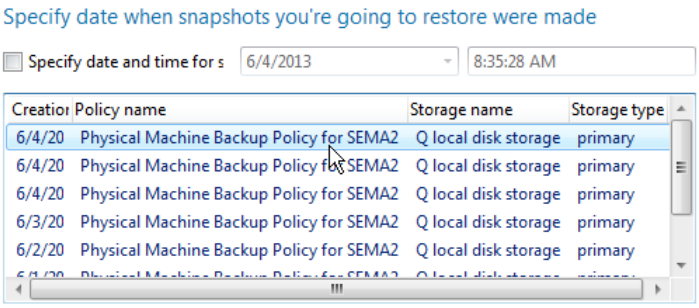

Despite the fact that you're allowed to initiate complete restore or retrieval of certain files/folders from invalid backup images, please do it at your own risk. Please consult the Administering storage backup data chapter to learn how to check images for integrity.

5. Click on **Restore selected volumes from backup**.

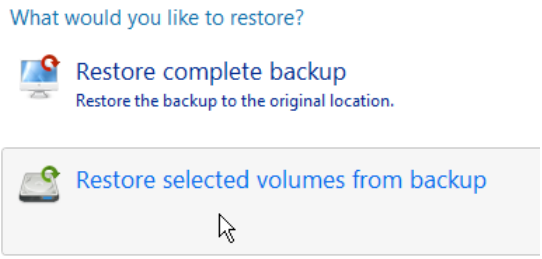

6. Specify data volumes to restore. If you choose system volumes, please additionally mark the corresponding option to use our WinPE recovery environment on-site, as it's the only option to restore system volumes of physical machines at the moment. Otherwise, the restore policy will fail. Click **Restore** to initiate the operation.

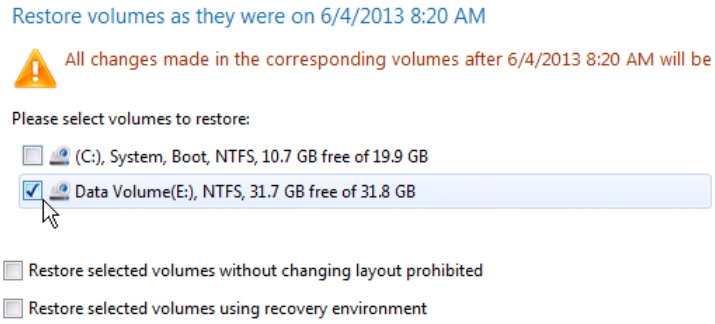

7. You will be informed on the operation start through a popup window.

8. Additionally we recommend you to create a notification sub-policy for easier monitoring.
9. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
10. When the restore task is over, its status will be updated.

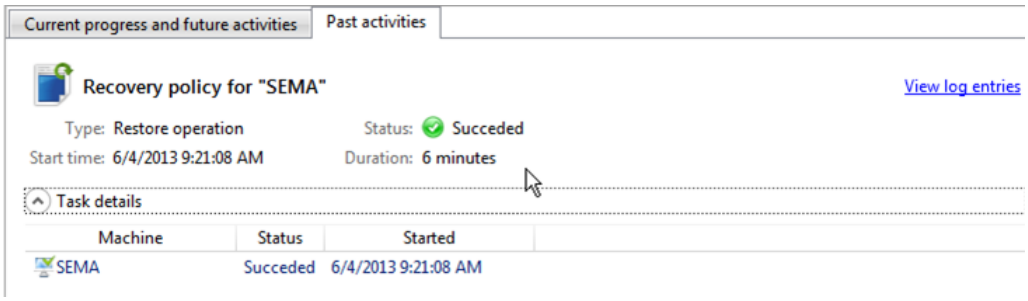

You can also accomplish restore of non-system volumes with the WinPE recovery media.

To know how to manage created policies, please consult the Managing Policies chapter.

Restoring an Entire Physical Machine or System Volumes by ID

Restore of an entire physical machine or system volumes involves the use of our WinPE recovery environment on-site as it's the only option to restore system (in-use) volumes of physical machines at the moment. You've got two options:

- Boot the target machine from the recovery media, connect to the infrastructure, browse backup storages for the required backup image, and finally initiate the restore operation. To know more on the subject, please consult the Configuring Recovery Policy from the WinPE Environment scenario;
- Create a restore policy in Console, and then send the generated helpdesk ID to the user. The user should only start up the failed computer from the recovery media and enter the obtained ID. The restore operation will be accomplished automatically, thus avoiding any mistake from the user's side.

Let's see in details how to restore an entire physical machine in the semi-automatic mode:

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target machine should have a network connection to Administration Server and Backup Server.
- The target machine should have at least 4GB of RAM.
- The hard disk of the target machine should be identical in size to that of the restored item (entire disk or certain volumes) or larger.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a Physical Machine Backup**,

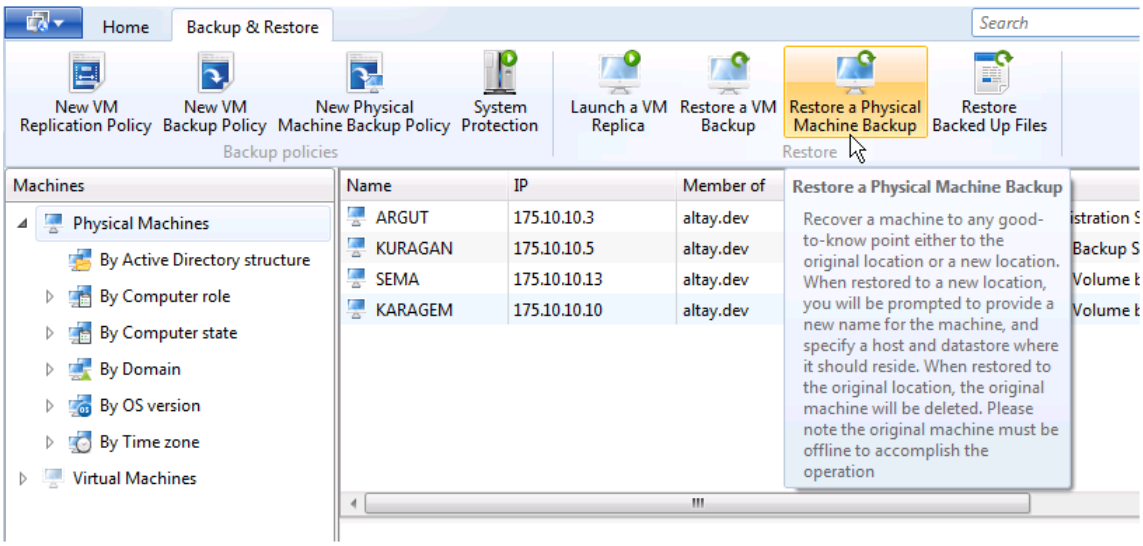

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a Physical Machine Backup**.

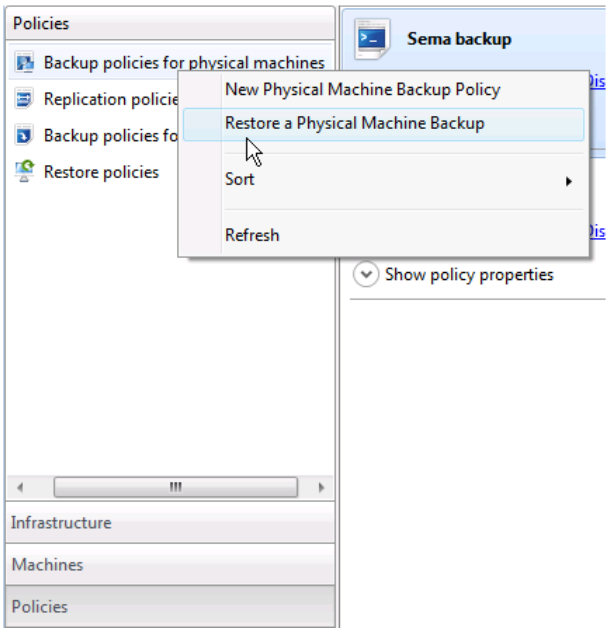

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name.

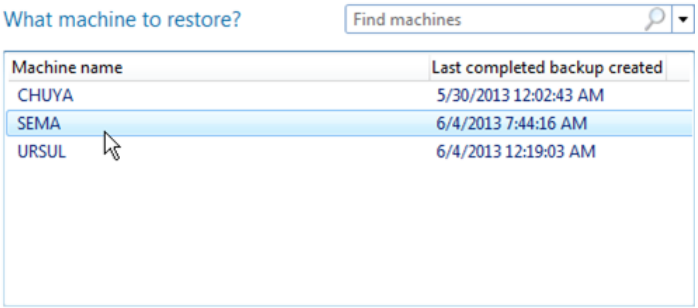

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

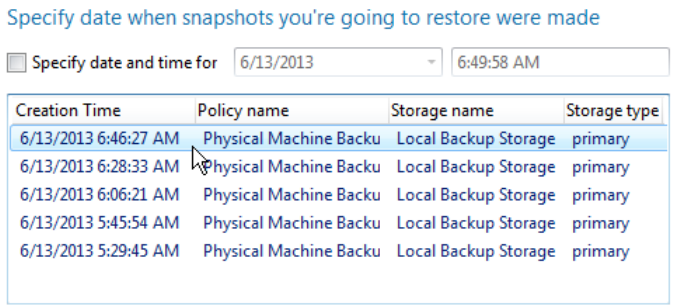

Despite the fact that you're allowed to initiate complete restore or retrieval of certain files/folders from invalid backup images, please do it at your own risk. Please consult the Administering storage backup data chapter to learn how to check images for integrity.

5. Click on **Restore complete backup** or **Restore selected volumes from backup** depending on your task. If you're going to restore data volumes only, please consult the Restoring Non-system Volumes Remotely scenario.

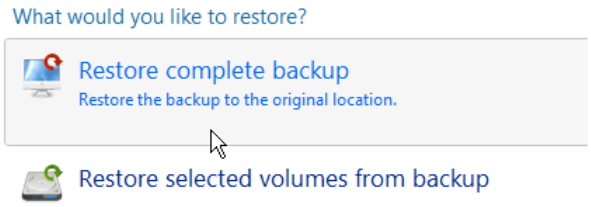

6. Either use the proposed helpdesk ID, or enter your own in the corresponding field. Click **Create policy** when ready.

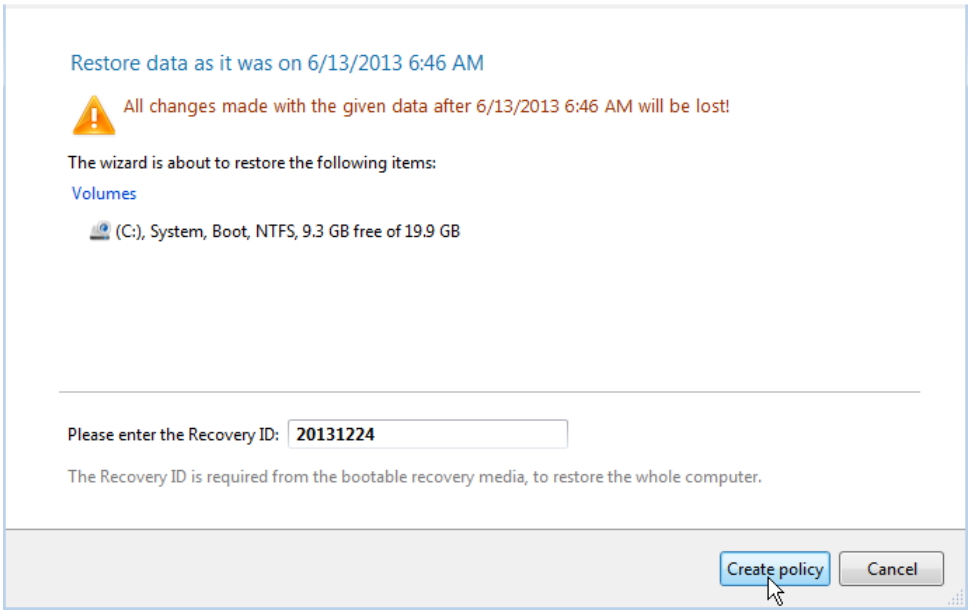

7. Save somewhere the resulted helpdesk ID. You can do it later by going to **Policies >** right click on **Restore policies**, then selecting the required policy and using the copy/paste function.

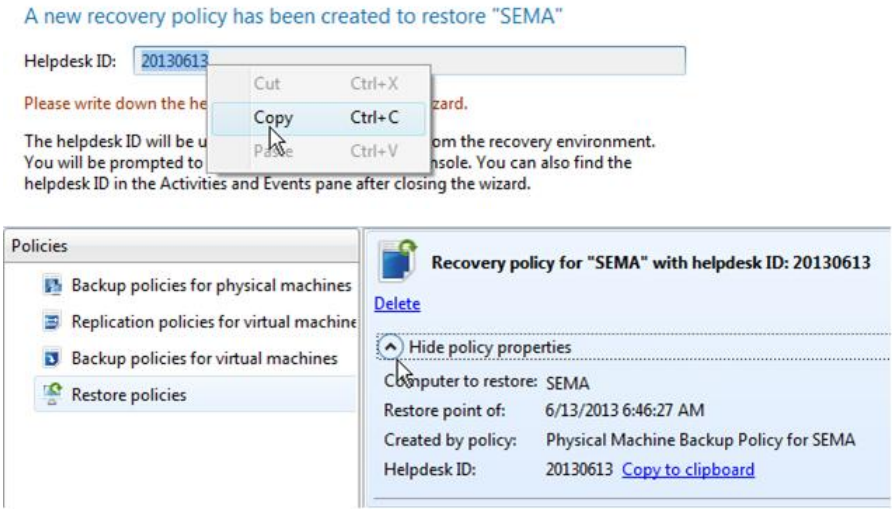

8. Boot the target machine from the WinPE recovery media, or send it to someone close to the target machine to let him/her initiate the restore operation.
9. Launch the **Bare Metal Recovery** Wizard.

10. Go through the wizard's welcome page, then select **Connect and use existing recovery policy**.

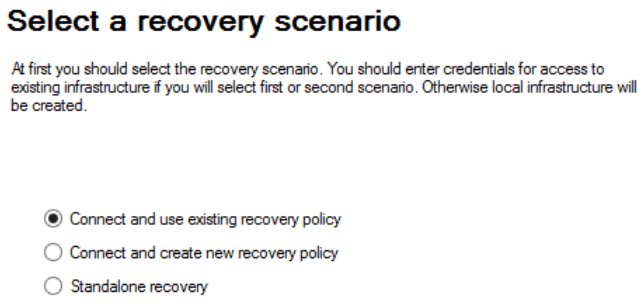

11. Provide a DNS name or IP address of Administration Server and its access credentials. Once done, the wizard will attempt to connect to Administration Server.

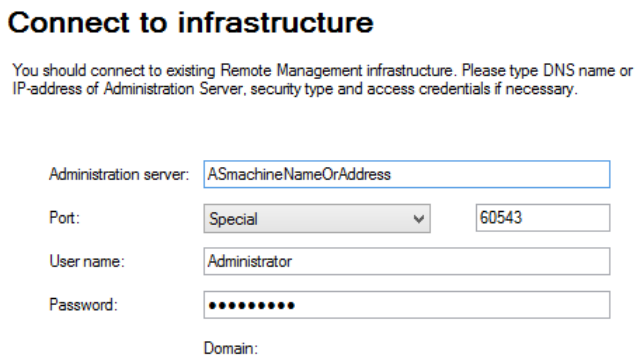

If a success, the wizard will proceed to the next step. If the wizard has failed to connect to the infrastructure, you can try to take a number of actions:

- Please check you have entered a correct DNS name or IP address of Administration Server.
- Close the wizard and run Network Configurator to check if there's access to network resources. If not, probably you should inject a network card driver.
- If there are several network cards on the machine, select the one that shares the same network with the target machine.
- Ping the target machine to make sure it's online.
- Please check the PRM service is running on the target machine.

12. Enter the saved earlier recovery policy ID in the corresponding field.

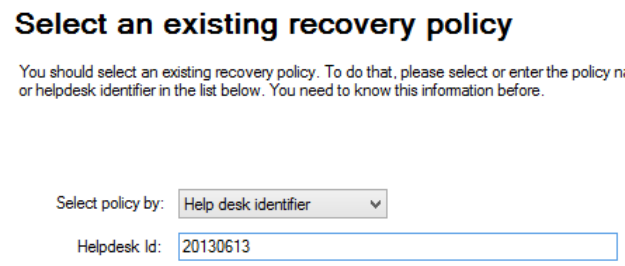

You can also find a desired recovery policy by name from the list of available policies. If a recovery policy has been created, but is not present in the list, press the **Refresh** button at the bottom of the list to update the information. Press **Apply** button to initiate the restore operation.

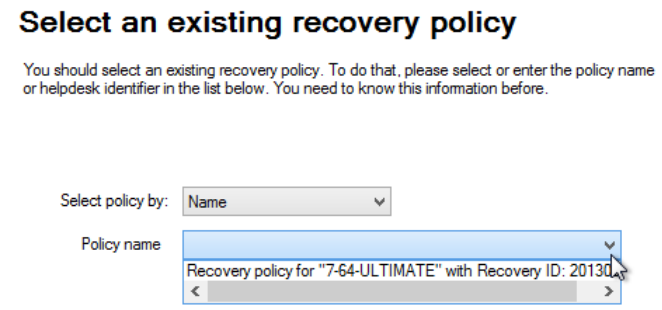

A restore operation can only be cancelled during data writing, but not directory and credentials replication.

13. The restore operation can be monitored in one of the consoles (the main GUI console, or PowerShell console).

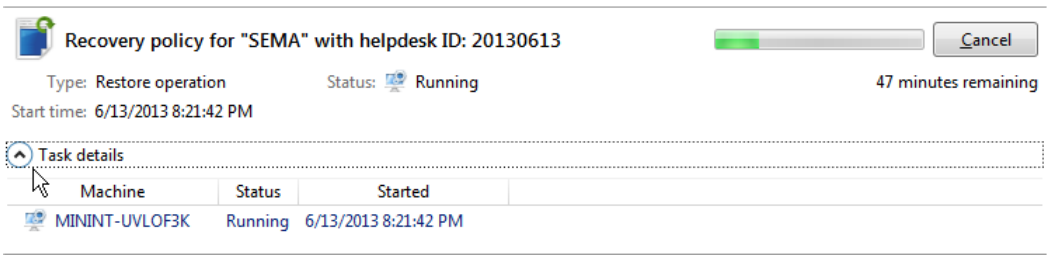

Once a restore operation by ID has been initiated, the corresponding restore policy will be deleted.

14. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
15. When the restore task is over, you can see its summary on the corresponding page.

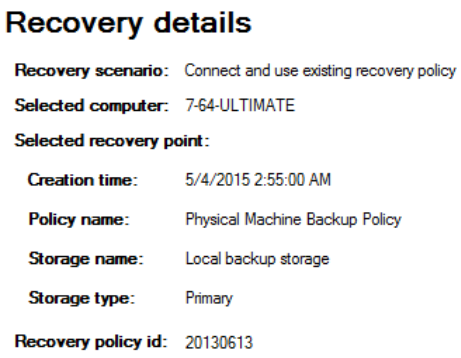

Configuring Recovery Policy from the WinPE Environment

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target machine should have a network connection to Administration Server and Backup Server.
- The target machine should have at least 4GB of RAM.

Operation scenario

1. Boot the target machine from the WinPE recovery media.
2. Launch the **Bare Metal Recovery** Wizard.

3. Go through the wizard's welcome page, then select **Connect and create new recovery policy**.

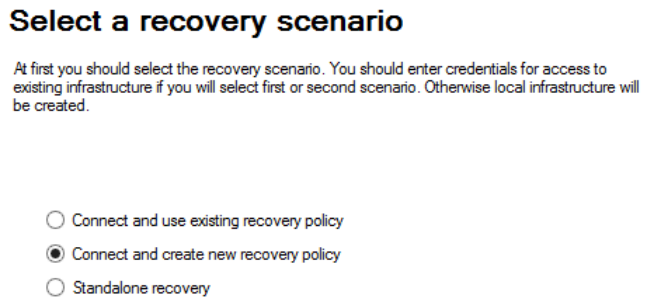

4. Provide a DNS name or IP address of Administration Server and its access credentials. Once done, the wizard will attempt to connect to Administration Server. If a success, wizard will proceed to the next step. If not, please make sure the target machine is available and the PRM service is running on it – click here to know more on the subject.

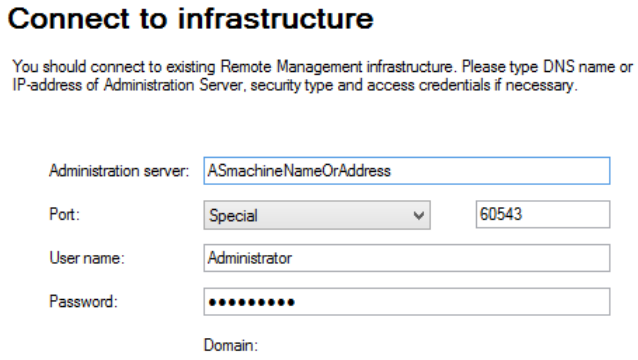

5. Select one of the backed up earlier physical machines. This list shows all the protected machines from all the backup storages of the PPR infrastructure.

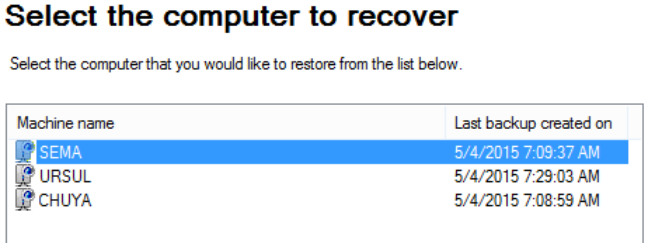

Machines, which backup images are controlled by an inaccessible or broken Backup Server are not shown in this list.

6. Then you need to choose a desired restore point, if several. The latest (at the top) recovery point is selected by default. Sessions from both secondary and primary storages are shown in this list. To hide sessions from secondary storages, set the **Only show backups residing on primary storages** option.

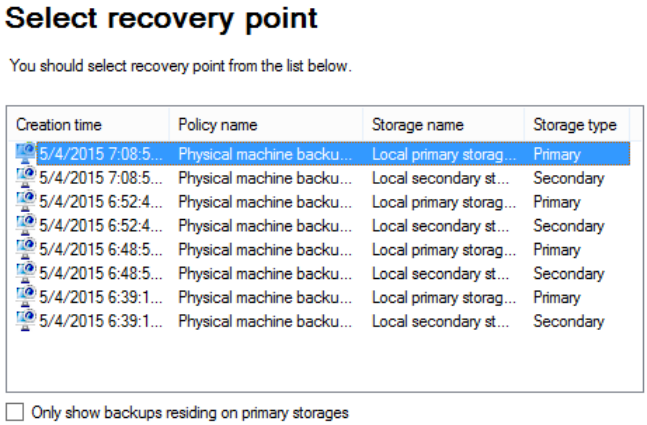

File-level backups and MS Exchange backups are invisible in this session list.

Invalid sessions, which didn`t pass integrity checkup, are shown with warning icon and they are disabled for selection.

7. Use one of the following scenarios to configure the recovery policy:
   - Restoring a Single Non-system Volume to Unallocated Space;
   - Restoring an Entire Machine or System Volumes to Original Location;
   - Bare-metal Recovery to Dissimilar Hardware.
8. Once the policy is configured and started, the restore operation can be monitored in one of the consoles (the main GUI console, or PowerShell console).
9. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.
10. When the restore task is over, its status will be updated.

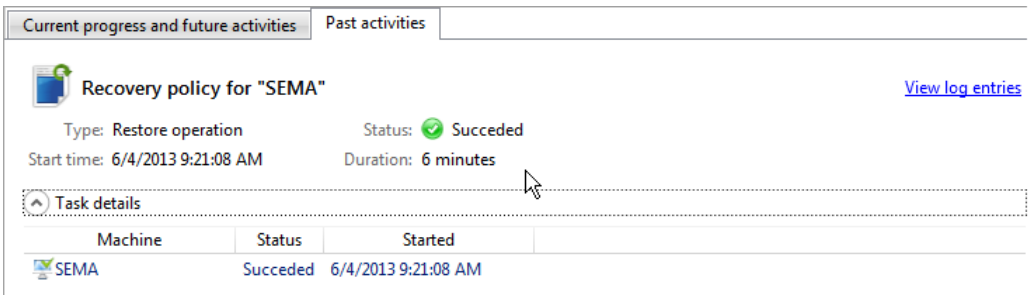

Standalone Recovery with no Connection to the PPR Infrastructure

This scenario can help to recover system or data volumes of a physical machine without connecting to the PPR Infrastructure. This scenario is the only option when Administration Server and/or Backup Server are down.

The entire restore procedure is accomplished from the WinPE recovery media. You can attach and use backup storages from:

- A local disk of the target machine;
- An external storage device connected to the target machine;
- A network share.

No synchronization with main PPR infrastructure is required in the standalone restore mode, as it involves creation of a special temporary PPR infrastructure.

Once the restore wizard is closed, the temporary infrastructure and a corresponding recovery policy will be deleted. If you start the wizard once again, all log files from the previous restore operation will be overwritten.

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target machine should have access to the required backup storage (local or network).
- The target machine should have at least 4GB of RAM.
- Network shares located on a machine under control of a non-server version of Windows OS are not supported. Use shares of Windows Server or Linux.

Operation scenario

1. Boot the target machine from the WinPE recovery media.
2. Launch the **Bare Metal Recovery** Wizard.

3. Go through the wizard's welcome page, then select **Standalone recovery**.

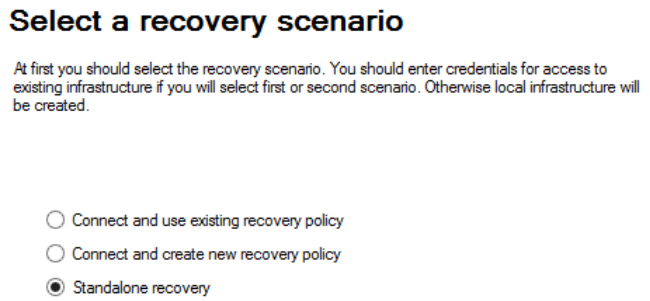

4. Set a path to the required backup storage (local, external, network). It can be any type of storage (primary, secondary, deduplicated or not):

   Storage on a local disk or external storage device:

   - Select **Local** as the preferred storage type. Enter a path to a folder with storage or select it using Browse button.

WinPE reassigns drive letters automatically, so they can differ from letters in the original system.

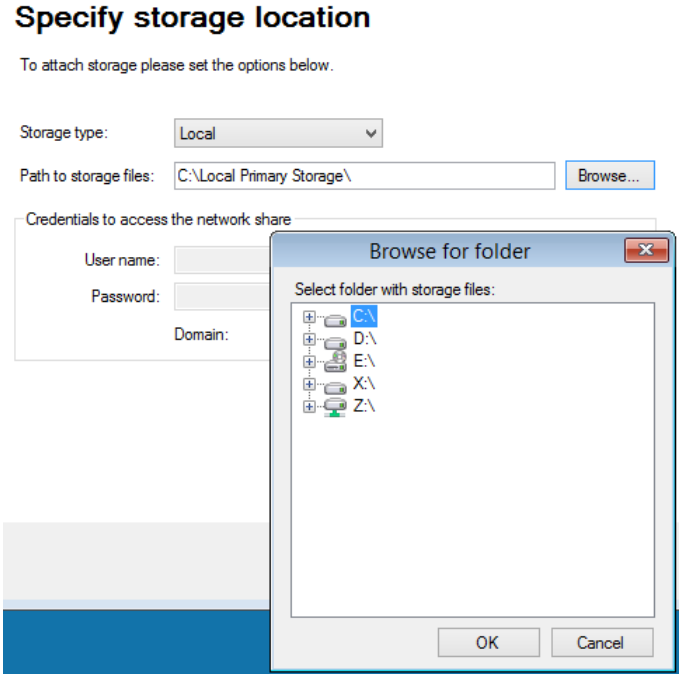

Storage on a network share:

- Select **Network** as the preferred storage type. Enter a network path and access credentials.

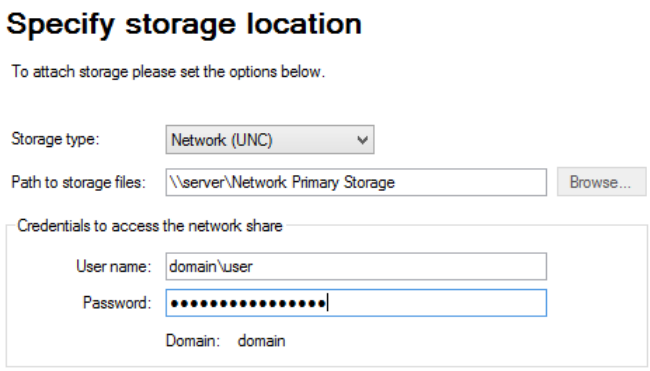

If there's no PPR storage by the provided path, or incorrect network credentials have been entered, the wizard will output a corresponding error.

5. If there has been found only one backup storage by the provided path, it will be selected by default. If there are several storages in the specified location, the wizard informs you about it, prompting to choose one of them to proceed. To help you make the right choice it also outputs a number of storage properties at this stage. Select the desired storage, then click **Next** to initiate the attachment.

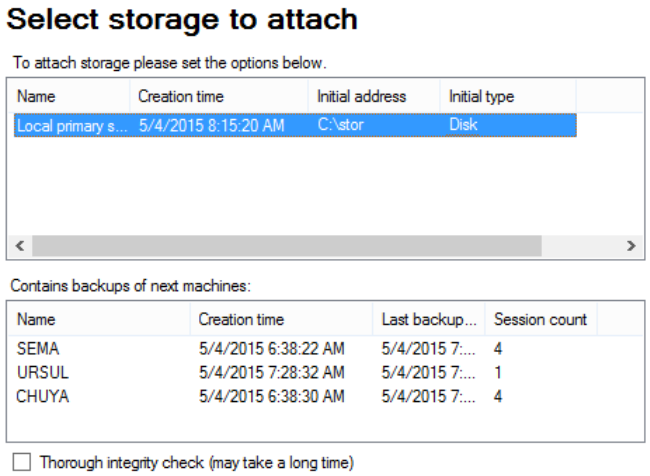

During the attachment process, backup data will be checked for integrity. If you'd like to additionally check CRC of backup data, please mark the **Thorough integrity check** option. Please note that verification of CRC requires significantly more time to complete.

6. If the specified storage contains deduplicated backup sessions, you will be prompted to select an appropriate Deduplication Server.

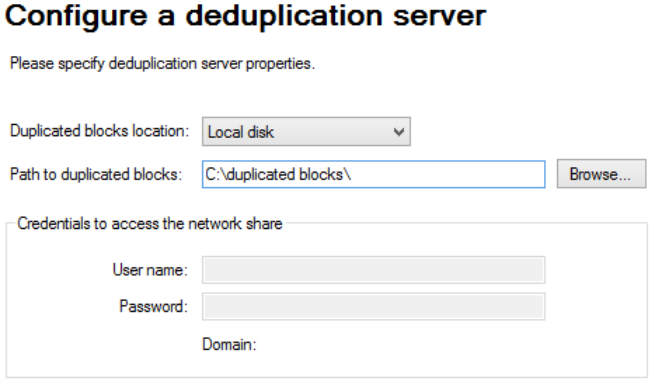

If there are no deduplicated blocks by the provided path that belong to the specified storage, the wizard will output an error.

7. As soon as the storage attachment is completed you need to choose a desired restore point, if several. The latest (at the top) recovery point is selected by default.

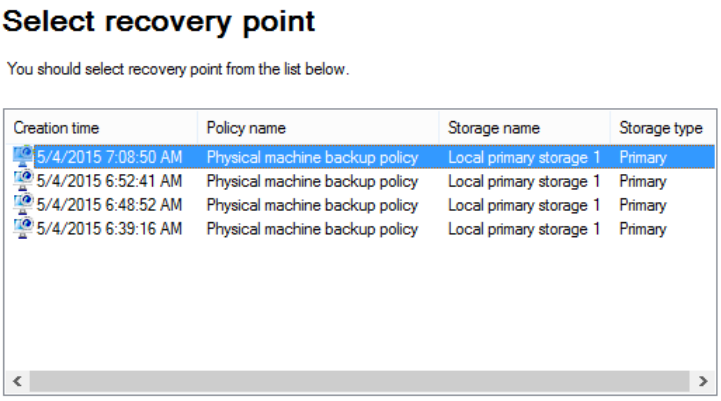

File-level backups and MS Exchange backups are invisible in this session list.

Invalid sessions, which didn`t pass integrity checkup, are shown with warning icon and they are disabled for selection.

8. Use one of the following scenarios to configure the recovery policy:
   - Restoring a Single Non-system Volume to Unallocated Space;
   - Restoring an Entire Machine or System Volumes to Original Location;
   - Bare-metal Recovery to Dissimilar Hardware.

Restoring a Single Non-system Volume to Unallocated Space

This scenario requires the use of the WinPE recovery media. You're allowed to restore any volume from a selected backup session to a block of unallocated space on the target machine hard disk. Existing partitions and data on the target hard disk won`t be affected by this operation.

Despite the fact that any volume is allowed to restore to unallocated space, it is not recommended to use this mode for system volumes, as we cannot guarantee startup of an operating system on the target hard disk.

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target disk should contain one or several free blocks, that are not allocated by any partition. It can be an empty hard disk as well.

- If the target hard disk is MBR-type, please make sure it has at least one vacant primary slot. If it already has 4 primary partitions, or 3 primary plus 1 extended partition, the restore wizard will not be able to create one more primary partition.

Operation scenario

1. Start up the target computer from the prepared WinPE media.
2. Choose how you'd like to connect to the PPR infrastructure and select a machine and one of the existing backup sessions to restore according to one of the following scenarios:
   - Configuring Recovery Policy from the WinPE Environment;
   - Standalone Recovery with no Connection to the PPR Infrastructure.
3. Once the recovery point is selected according to one of the above scenarios, the wizard prompts to select the required restore mode (**Restore single partition…** is what we need).

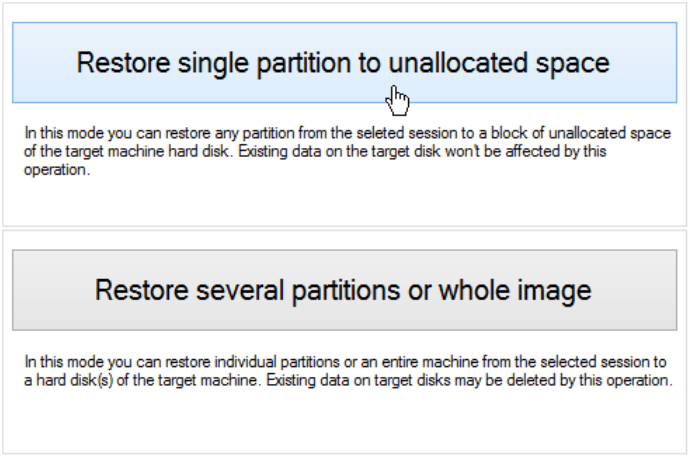

4. Specify a volume to restore and a free block where you'd like it to restore from the list of available in the system. By default, the restored volume will allocate the entire block (a 500GB volume restored to a 750GB free block will become 750GB when the operation is over). If you'd like to keep the original volume size, please unmark the **Resize volume(s) proportionally** option.

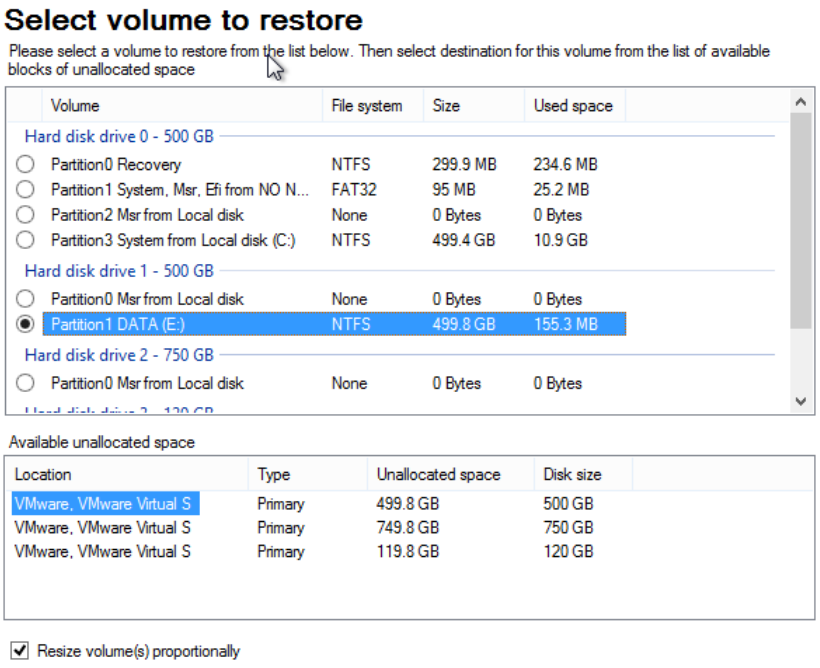

Restore to a free block that is smaller in size than the original volume is only possible when the volume resize is allowed.

Resize of MS Reserved, MS Recovery and ExFAT volumes is not allowed, thus the corresponding option will be shadowed for any of the mentioned objects.

5. Review all introduced changes, and then confirm the operation. A volume selected for restore is marked with a red frame.

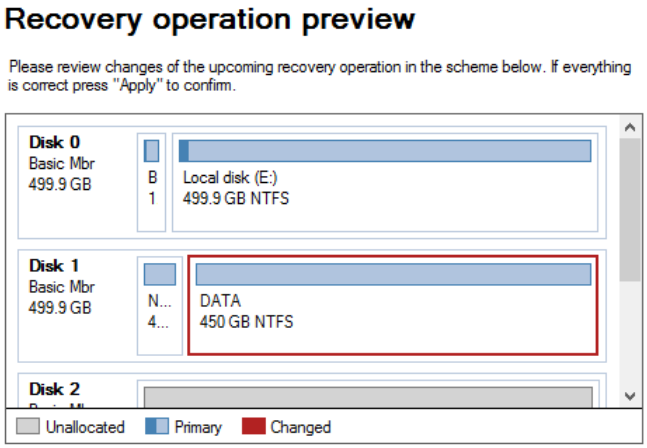

6. When the restore task is over, you can see its summary on the corresponding page.

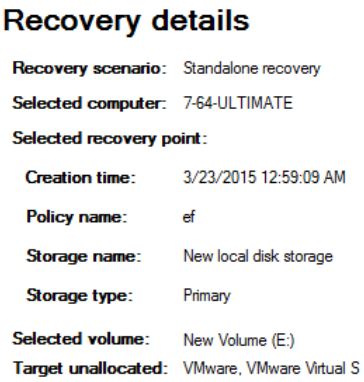

Restoring an Entire Machine or System Volumes to Original Location

This scenario requires the use of the WinPE recovery media. We recommend it to you when OS and/or file system have been damaged to get everything back on track with minimal effort (no additional boot correction or OS adjustment are needed).

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target machine should have one or several hard disks that are identical in size and layout with those stored in a backup image.

Operation scenario

1. Start up the target computer from the prepared WinPE media.
2. Choose how you'd like to connect to the PPR infrastructure and select a machine and one of the existing backup sessions to restore according to one of the following scenarios:

- Configuring Recovery Policy from the WinPE Environment;
- Standalone Recovery with no Connection to the PPR Infrastructure.

3. Once the recovery point is selected according to one of the above scenarios, the wizard prompts to select the required restore mode (**Restore several partitions…** is what we need).

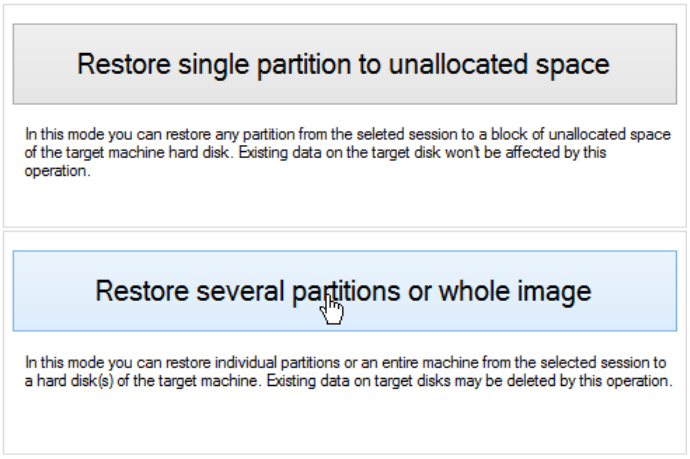

4. By default the wizard selects all volumes from the specified backup image and displays the detected original target hard disk(s) in the 'Available destination' section.

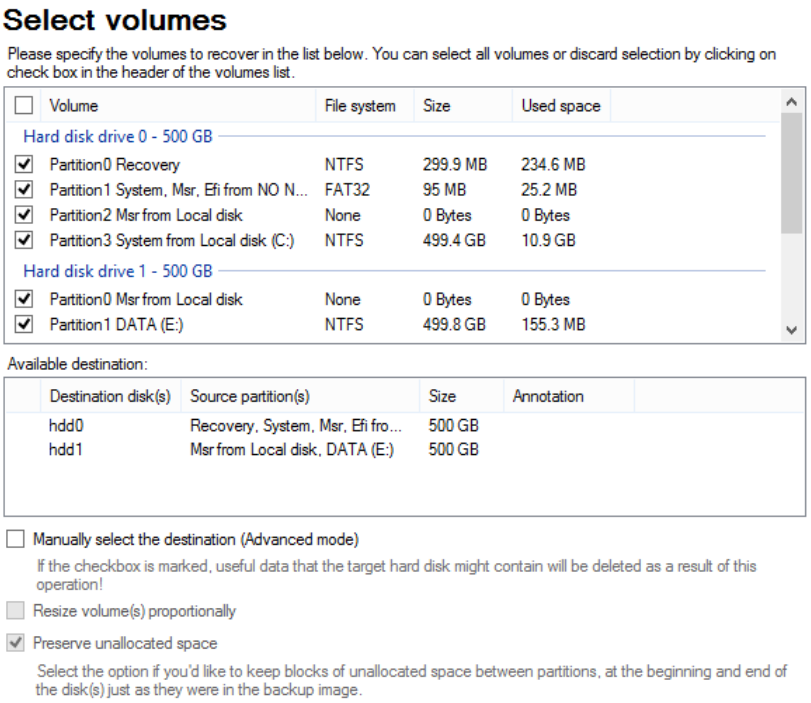

If the wizard has failed to find one or several disks which are planned for restore, it will output a corresponding error.

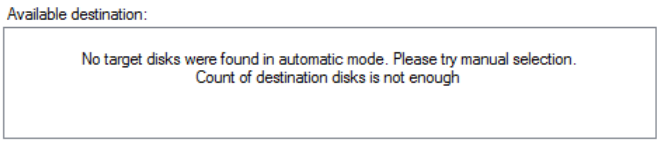

You're free to deselect any volume you do not want to be recovered and thus rewritten. Please note that use of different time stamps for different volumes may lead to data inconsistency.

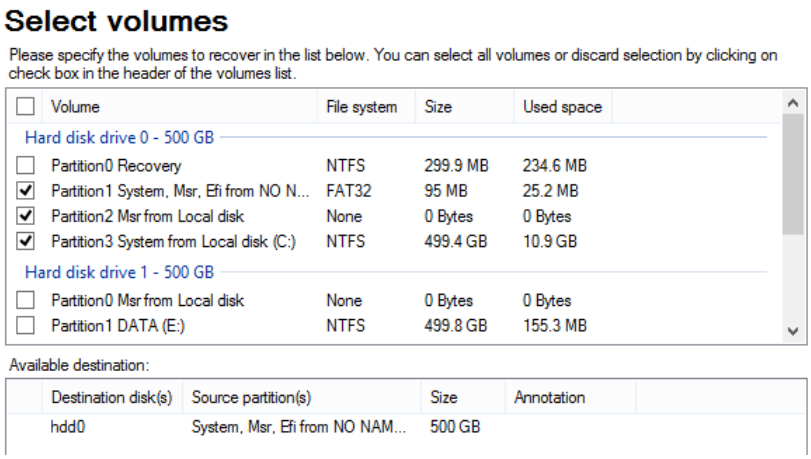

It is highly recommended not to deselect dependent system volumes, if they are present in a backup image. If you do this, OS might not start up. When trying to deselect a dependent system volume the wizard will warn you about it.

If you want to restore data volumes only, please deselect ALL system volumes.

5. Review all introduced changes, and then confirm the operation. Volumes selected for restore are marked with red frames.

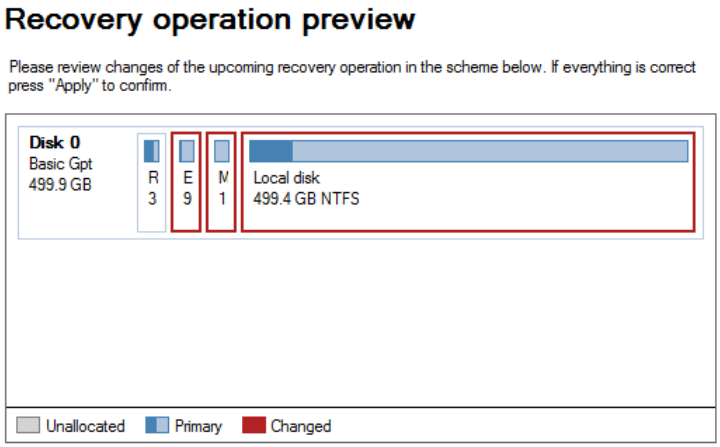

6. When the restore task is over, you can see its summary on the corresponding page.

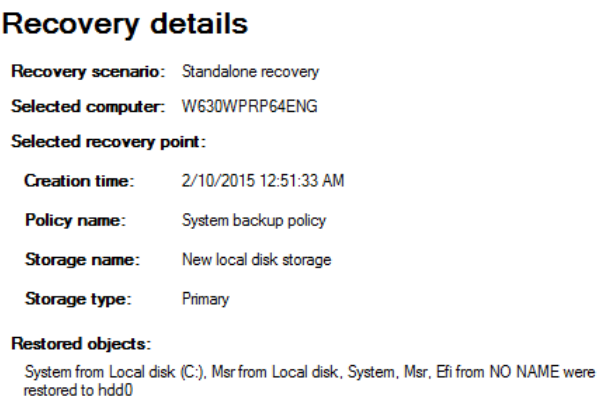

Bare-metal Recovery to Dissimilar Hardware

Let's assume one of the protected by PPR physical servers has failed due to a hardware failure. The failed hardware platform is quite obsolete and it's next to impossible to find and replace the damaged hardware devices. Disaster recovery to a new hardware platform seems the best way out. No problem – do a bare-metal recovery with the P2P option to guarantee Windows OS (any since Windows XP) will start up on the new hardware.

Prerequisites

- You should have a WinPE recovery media prepared with Recovery Media Builder.
- The target machine should have a network connection to Administration Server.
- The target machine should have at least 4GB of RAM.

Operation scenario

1. Start up the target computer from the prepared WinPE media.
2. Choose how you'd like to connect to the PPR infrastructure and select a machine and one of the existing backup sessions to restore according to one of the following scenarios:
   - Configuring Recovery Policy from the WinPE Environment;
   - Standalone Recovery with no Connection to the PPR Infrastructure.
3. Once the recovery point is selected according to one of the above scenarios, the wizard prompts to select the required restore mode (**Restore several partitions…** is what we need).

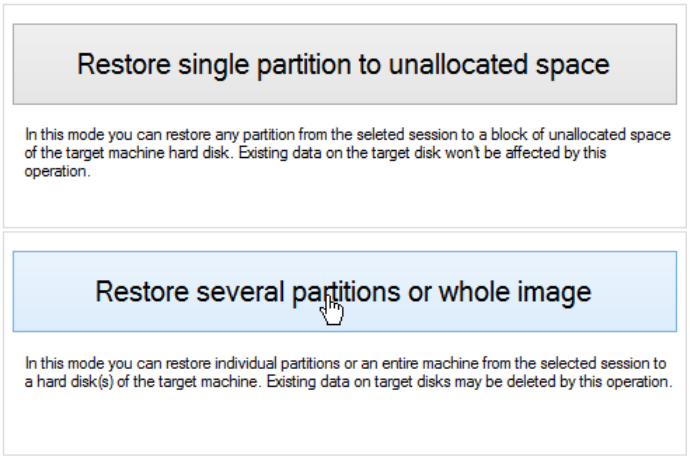

4. By default the wizard selects all volumes from the specified backup image and displays the detected original target hard disk(s) in the 'Available destination' section.

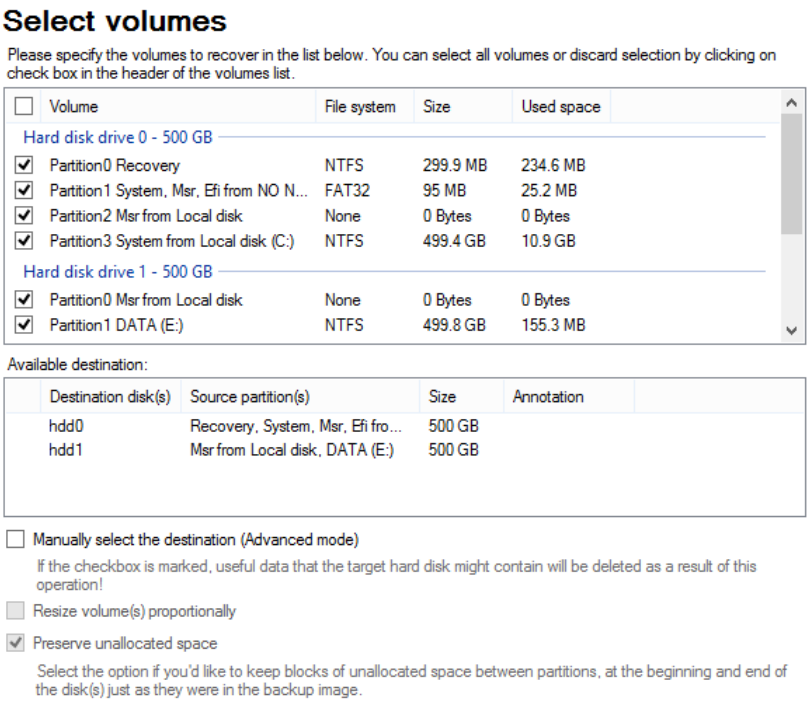

You're free to deselect any volume you do not want to be recovered and thus rewritten. Please note that use of different time stamps for different volumes may lead to data inconsistency.

It is highly recommended not to deselect dependent system volumes, if they are present in a backup image. If you do this, OS might not start up. When trying to deselect a dependent system volume the wizard will warn you about it.

If you want to restore data volumes only, please deselect ALL system volumes.

If an appropriate hard disk(s) has been found, click **Next** to proceed. If not, it might happen due to the following reasons:

- WinPE fails to detect some disks, most probably because of missed storage drivers.

  Solution: Add drivers through a corresponding dialog.

- Target disks are not empty (contain some partitions and data) and they ARE NOT the original disks, i.e. their contents are not stored in the specified backup image.

  Solution:

  - Clean up the target disks, if they do not contain any valuable data or move their contents to another storage, then restart the Recovery Wizard;
  - Connect new empty disks to the target machine;
  - Mark the **Manually select the destination…** option to allow data overwriting, then select the desired target disk from the list. Please note that all data on the disk selected in the advanced mode will be completely deleted.

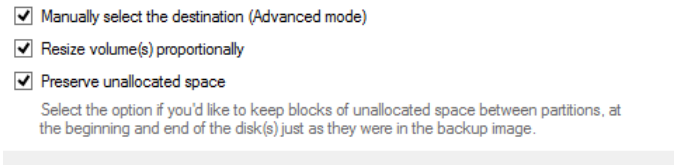

The advanced mode enables to select one target disk at a time.

- The target disk(s) is smaller than the original.

  Solution:

  - Deselect volumes from the backup image to fit in to the target configuration;
  - Connect a larger empty hard disk;
  - Mark the **Manually select the destination…** option to allow shrinking of volumes (resize down), then select the desired target disk from the list. The advanced mode allows proportional resize (up or down) of restored volume depending on the target disk size.

The advanced mode enables to select one target disk at a time.

5. Review all introduced changes, and then confirm the operation. Volumes selected for restore are marked with red frames.

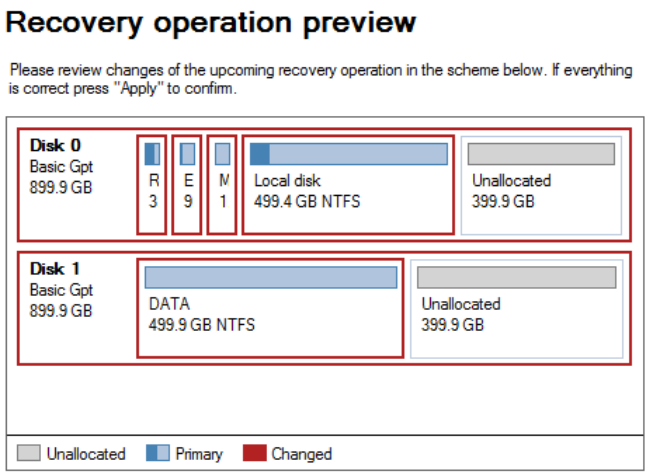

6. When the restore task is over, you can see its summary on the corresponding page. Click **Finish** to close the wizard and exit to the Express Launcher.

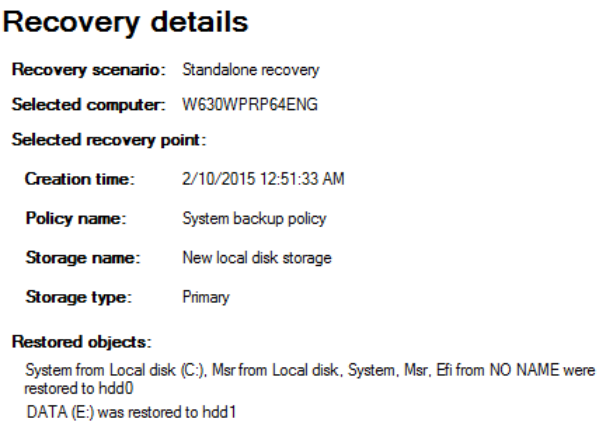

7. If it is needed, launch the **P2P Adjust OS Wizard** to adjust Windows OS to the new hardware. Before you start, please make sure you've got drivers for the new hardware ready to use, not zipped or in .exe files.

8. On the Wizard's Welcome page, click the **Next** button.
9. From the list of all found Windows systems (if several) select one you need to adjust to the new hardware. If you're willing to adjust them all, just re-launch this wizard for each.

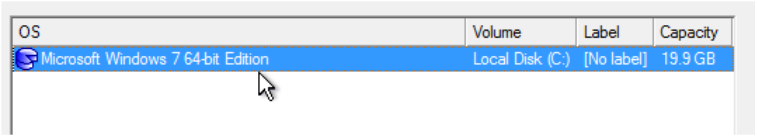

10. There are two execution modes to choose from: **fully automatic** and **advanced**. Below we will go set-by-step through the automatic scenario to show the whole process, and then take a closer look at specifics of the advance scenario. Select **Adjust the OS to the new hardware automatically**.

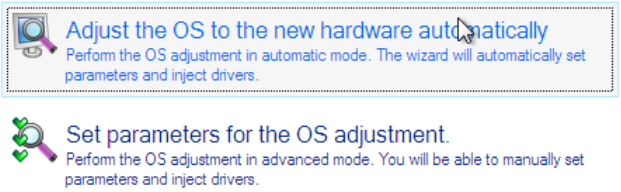

11. The wizard will automatically accomplish all the necessary actions.

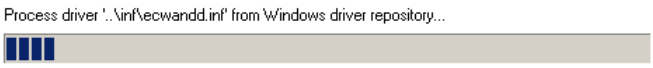

12. The only action that might be required from your side is to set a path to an additional driver repository in case the wizard has failed to find drivers for some boot critical devices in the built-in Windows repository. Generally together with new hardware you get its drivers for different operating systems on removable media (mostly CD or DVD). By collecting all these drivers in one folder you can let the wizard automatically pick and install only those required for your OS. Select **Search for drivers in a specific folder**.

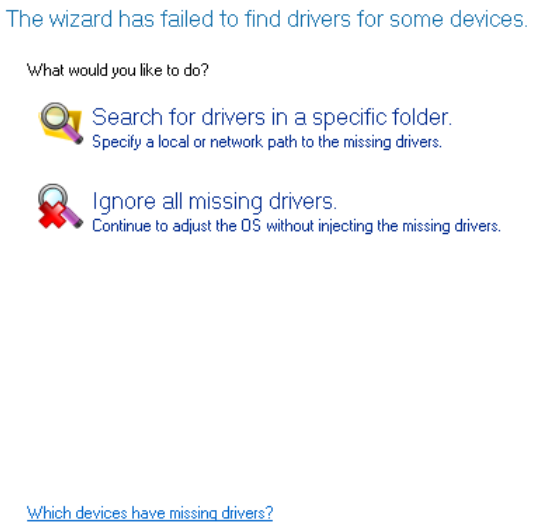

Click on the link at the bottom of the page to see what boot critical devices have no drivers. The wizard names all devices according to their model description, not some alphanumeric code, which is very convenient.

13. Though you've got the option to continue without injecting missing drivers for boot critical devices (The **Ignore all missing drivers** option), we strongly recommend you not to do it. Otherwise we cannot guarantee your Windows will start up on the new hardware.

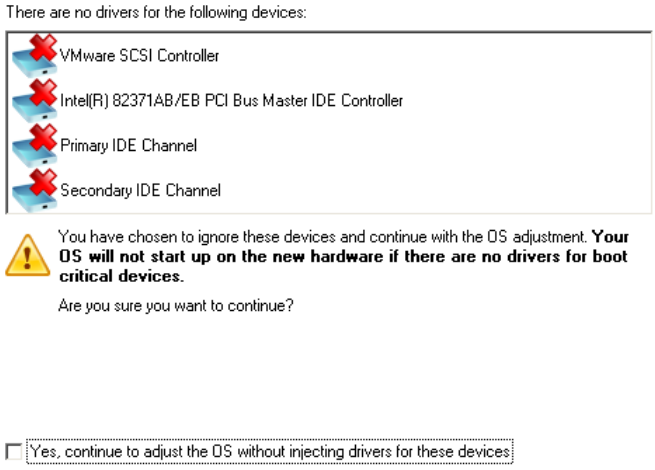

14. The wizard can search for drivers on a local disk or a mapped network share. In our case it's on a network share, this is why we need to map it first.

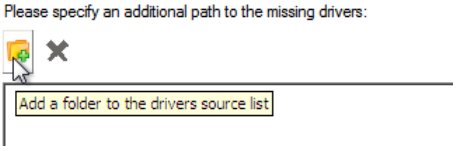

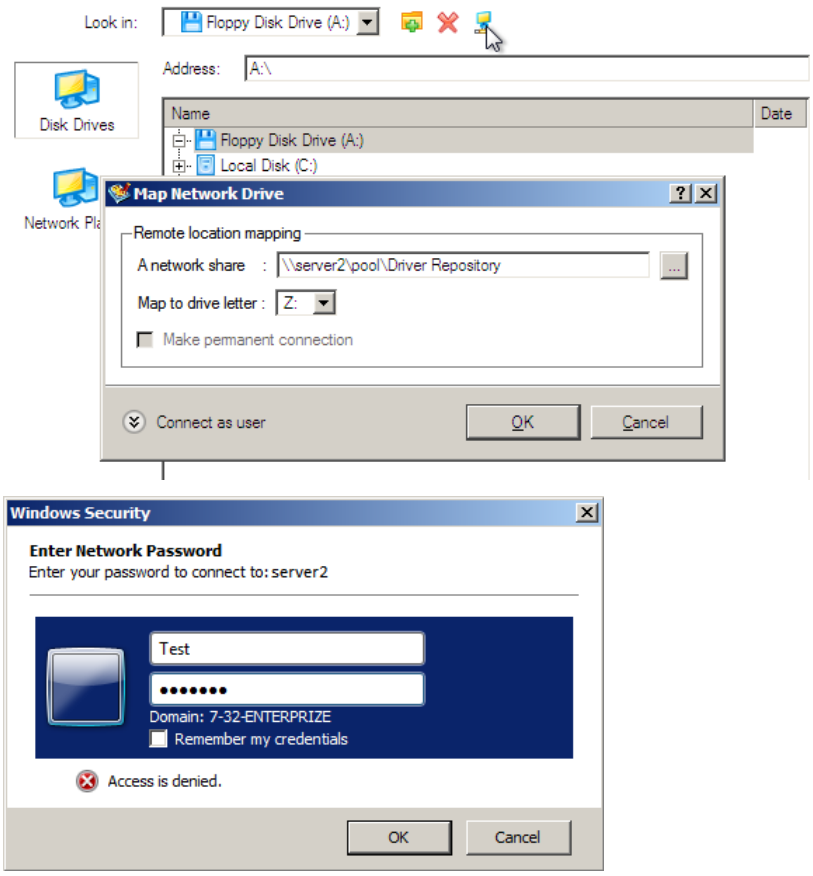

15. When done, we can select it as target.

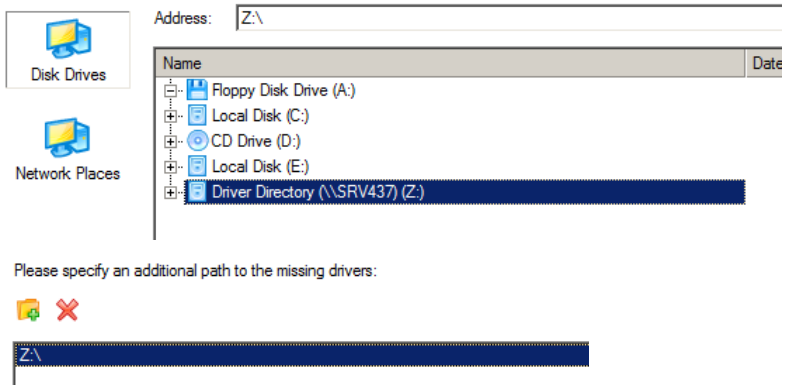

The wizard enables to specify several driver repositories.

16. If the wizard has found all missing drivers, it will ask you to confirm the operation. Apply the changes to complete.

After the operation is completed the system will be bootable on the new hardware. After the startup, Windows will initiate reconfiguration of all Plug'n'Play devices. It's a standard procedure, so please don't worry and prepare the latest drivers at this step to get the most out of the system.

Advance scenario specifics

1. To launch the advance mode, select **Set parameters for the OS adjustment**.

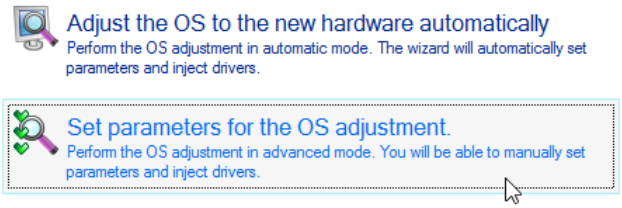

2. When setting additional driver repositories, you can specify how to process drivers for found hardware.

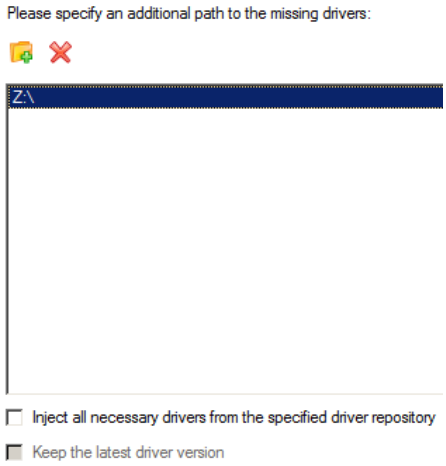

   - **Inject all necessary drivers…** Mark the checkbox to force injection of all drivers for your devices from the given driver repository(s), even if there are already installed drivers for some hardware. Please use this option if you suspect any of the installed drivers of not matching your hardware.
   - **Keep the latest driver version**. Mark the checkbox to keep the latest version of drivers during the forced re-injection. You can use this option only when the above option is active.

3. Just before the OS adjustment, you can additionally:

   - View all found hardware devices and their driver status by clicking . The wizard names all devices according to their model description, not some alphanumeric code, which is very convenient. So you can compare the listed devices with the given hardware to make sure the wizard has analyzed your system correctly.

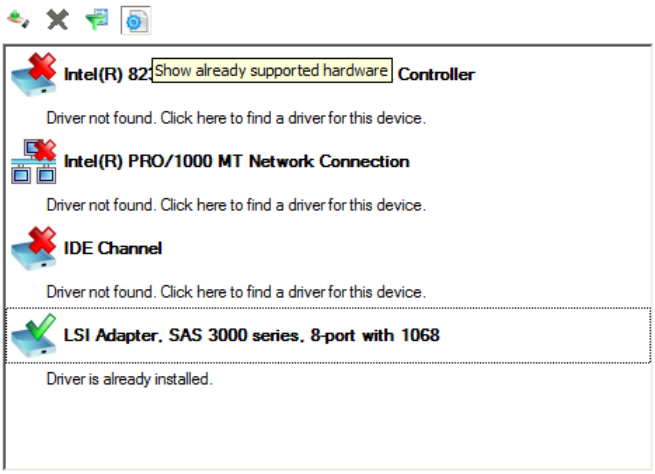

   - Filter devices without drivers by clicking . Unlike the automatic mode, where only boot critical devices (storage controllers) without drivers are being reported, here you can view and inject drivers for network cards as well.

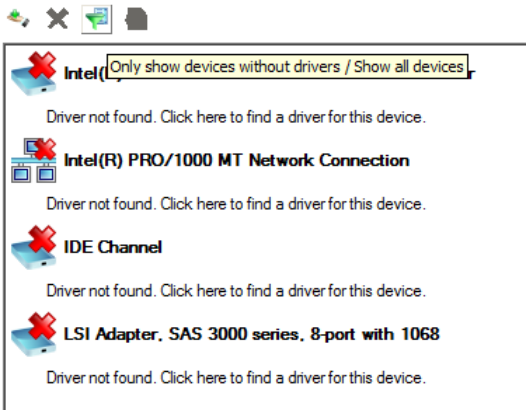

- Add a driver for each device that lacks it by clicking on the device, then browsing for the required location. The wizard will then match the device with drivers inside the given location and pick the right one.

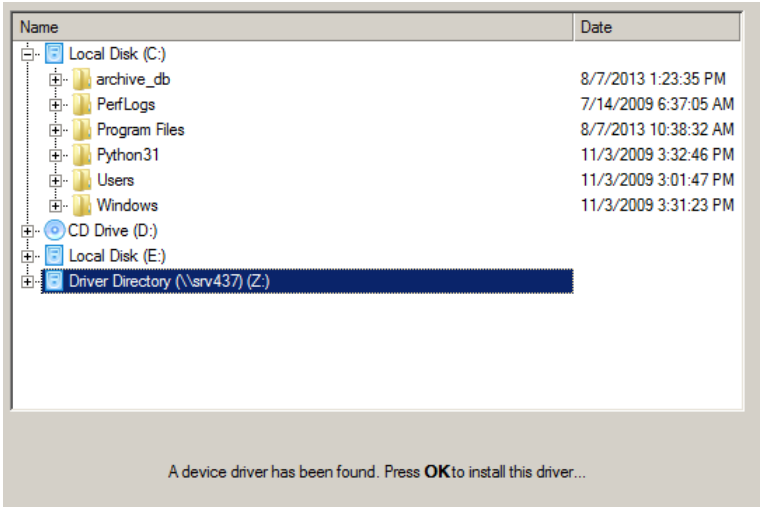

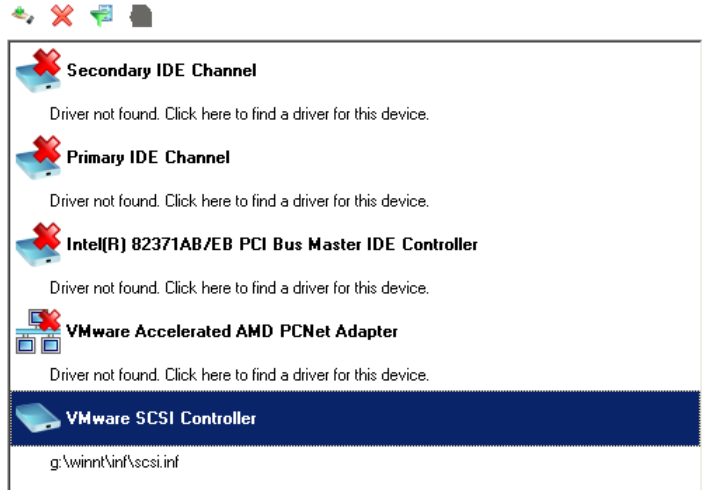

- Manually add a driver for a device that has not been found by our wizard by clicking , then specifying the required .INF file.

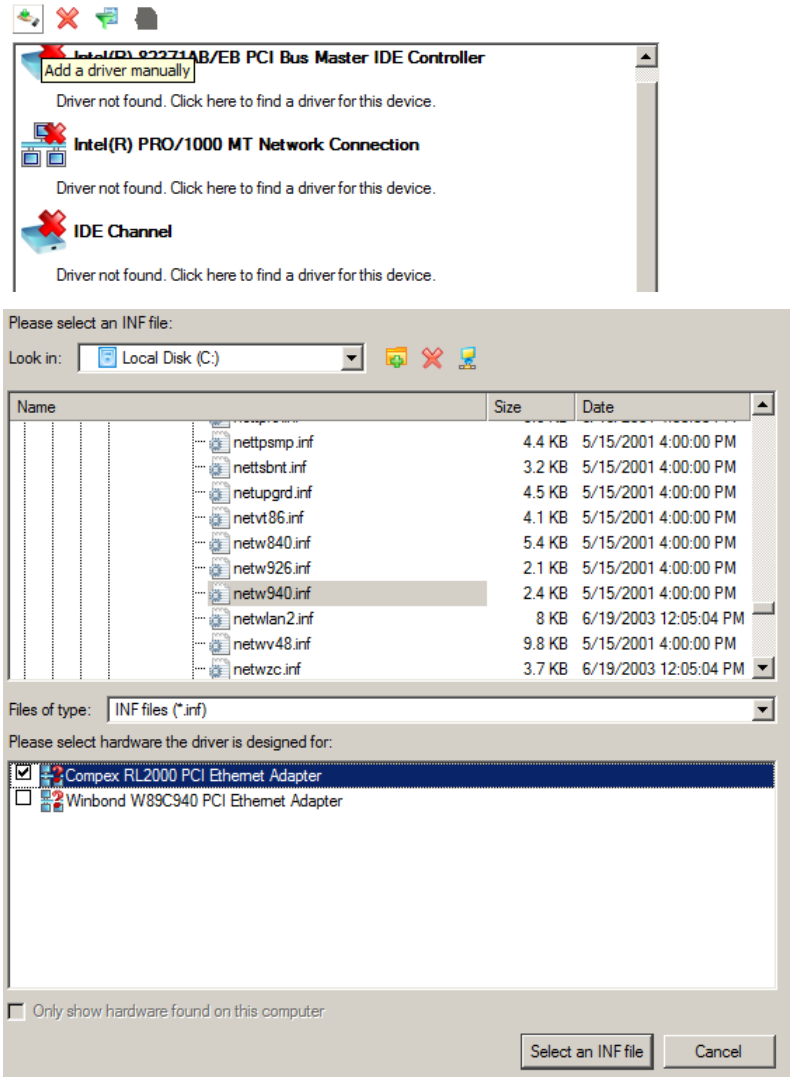

When selecting an .INF file that contains several driver records for hardware you both, have in the system and don't have, you can filter the list by marking the appropriate checkbox.

- Remove a driver for a device, which has not been found in the system.

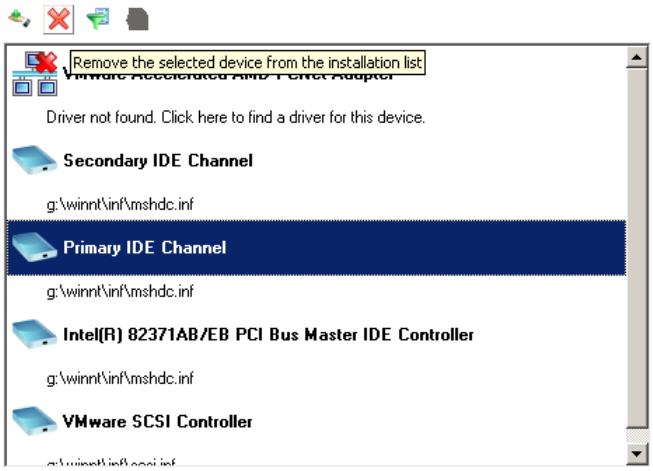

Launching Backup (Instant Restore)

The launch backup aka instant restore feature helps you minimize downtime of a failed production system. It enables to immediately run a Windows-based physical machine directly from one of available restore points in VMware ESX environment. Thus users may continue their activities, while you've got enough time to pinpoint and fix the failed system.

When applied to a physical system, PPR additionally does P2V migration to VMware ESX environment, which you can use for moving your production servers from physical to virtual hardware. To know more on the subject, please consult this scenario.

Using Tray Application

Prerequisites

To protect the host machine:

- Installed Backup Agent.
- Installed Tray Application.

To protect guest machines of the local Hyper-V host:

- Hyper-V that runs at least one virtual machine.
- Installed Backup Agent.
- Installed Tray Application.
- Installed Hyper-V Application.

Launching tray application

1. Launch Tray Application on the target machine by either clicking on its desktop icon or going to **Start > Programs > Paragon Protect & Restore > Protect & Restore Tray Application**.
2. Provide access credentials of a user that joins one of the PPR security groups and have enough privileges to accomplish desired tasks (monitoring, backup, backup and restore).

Monitoring backup activities

1. On the target machine please right click the Paragon's system tray application icon in the system tray, then select **Show activities**.

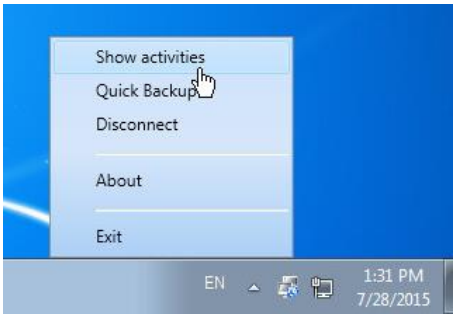

2. In the opened dialog you can see all backup policies that launched on the target machine (initiated by PPR Administrator or the user).

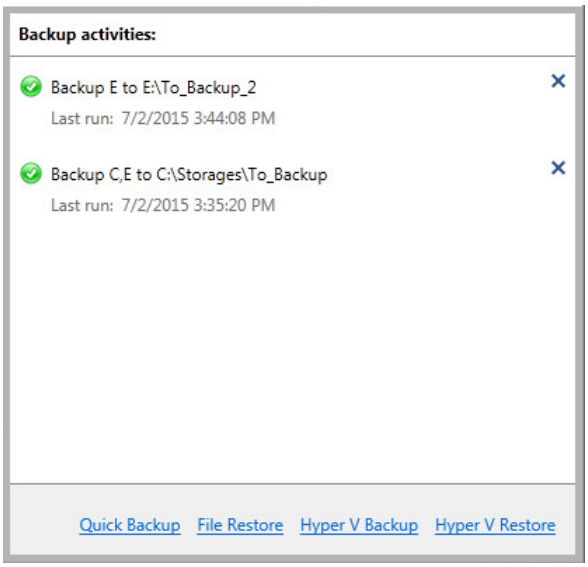

Hyper-V wizards become available only when all prerequisites are met.

Backing up the host machine

1. On the target machine please right click the Paragon's system tray application icon in the system tray, then select **Quick Backup**. You can also initiate creation of backup images through the Activities pane.

2. Specify a volume(s) you'd like to back up.

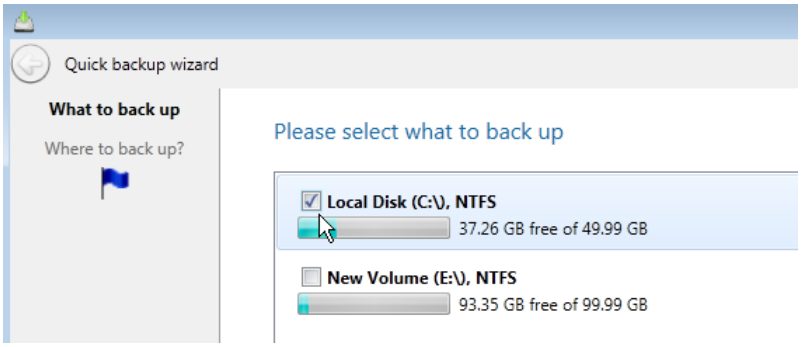

3. The wizard enables to save backup images either to a local folder, or a network share. Select the required backup location type in the corresponding menu.

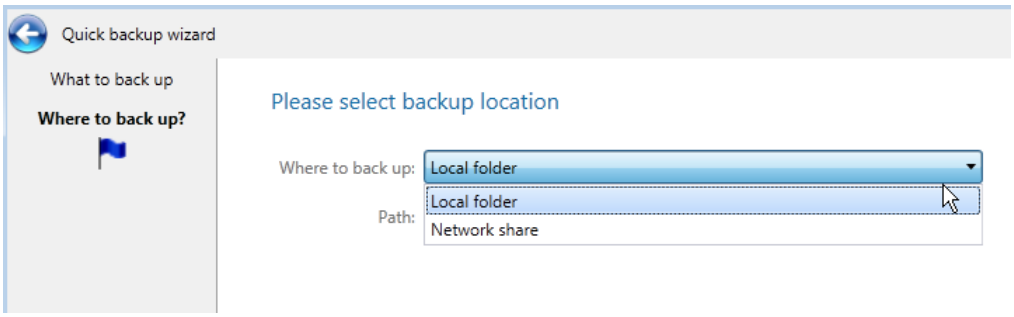

4. Depending on the selected backup location type, the wizard prompts you to specify a number of additional parameters:

For a local folder

- Click **Browse** to specify a local disk and folder to place backup images to. Use the **New folder** button if necessary. Please make sure the amount of free space on the selected volume is enough to store all planned backup images.

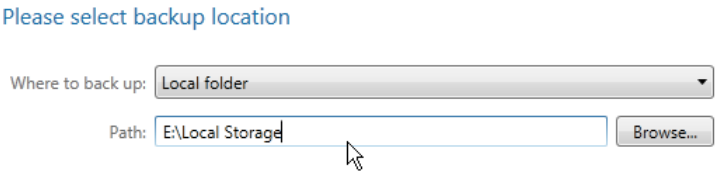

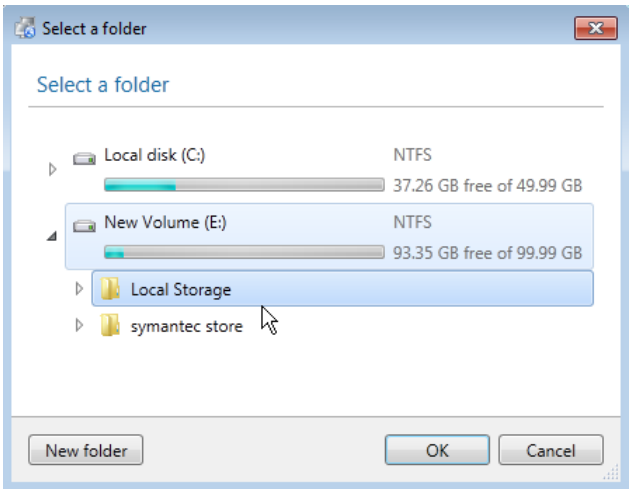

For a network share

- Provide a path to the required network share by manually entering its location or click **Browse** to find it on the net. Double click on the required network machine to get access to it. If necessary you will be prompted to provide login and password.

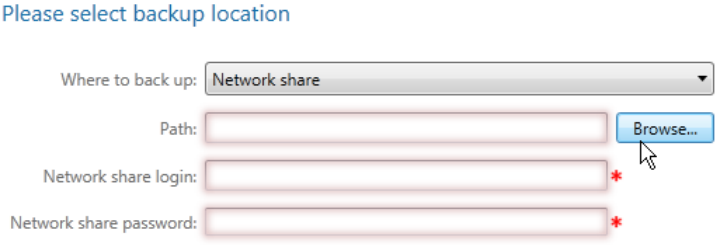

- If the provided credentials are valid, you will be able to browse the specified network machine for the required storage folder. Click **OK** when ready.

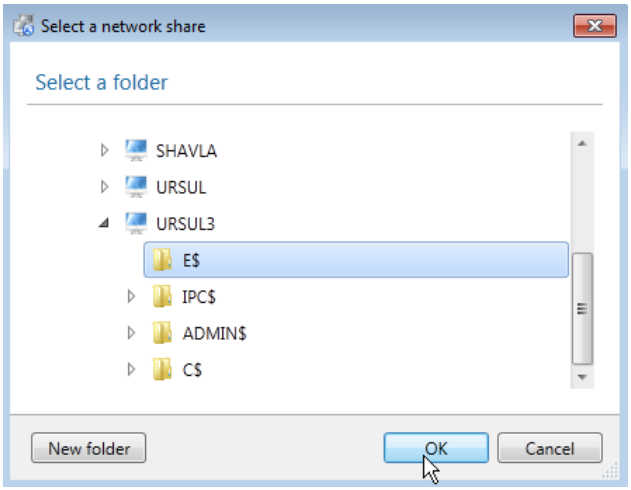

5. Click on the **Edit excludes** hyperlink to specify what data should be automatically ignored during backup. You can filter certain files or folders by creating masks.

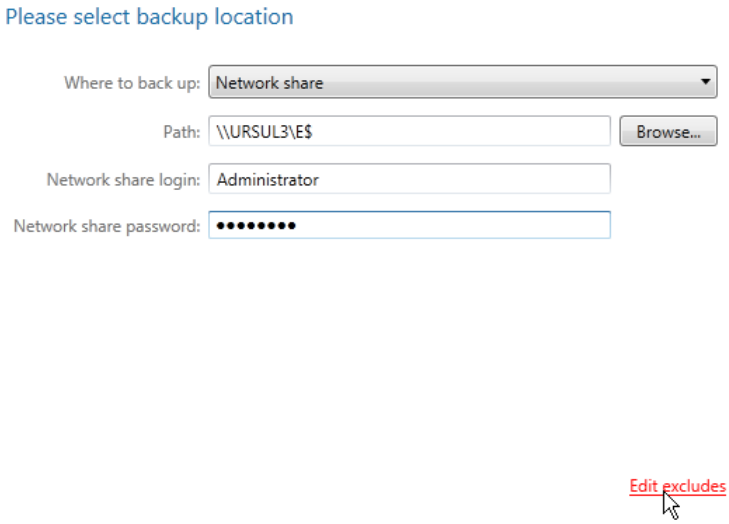

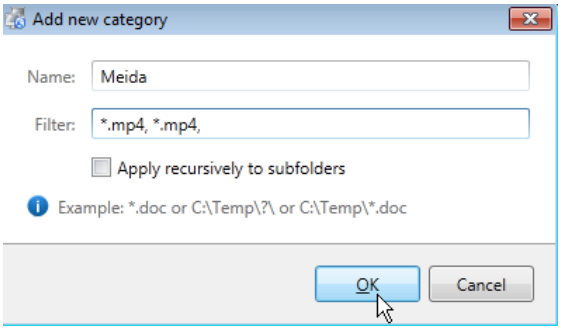

6. Click **Finish** when ready to initiate creation of a backup image.

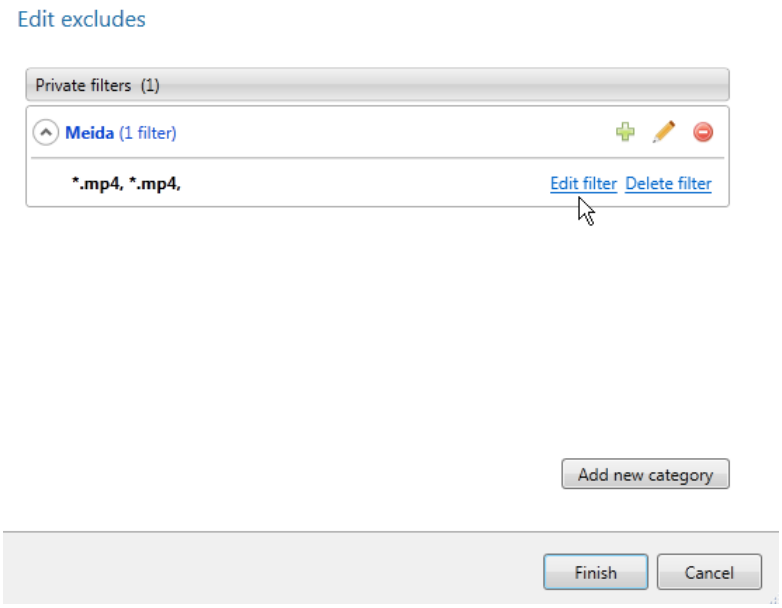

You will be informed on the operation start through a popup window.

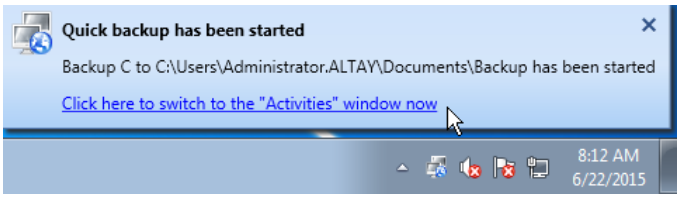

You can monitor the backup progress through the Activities pane.

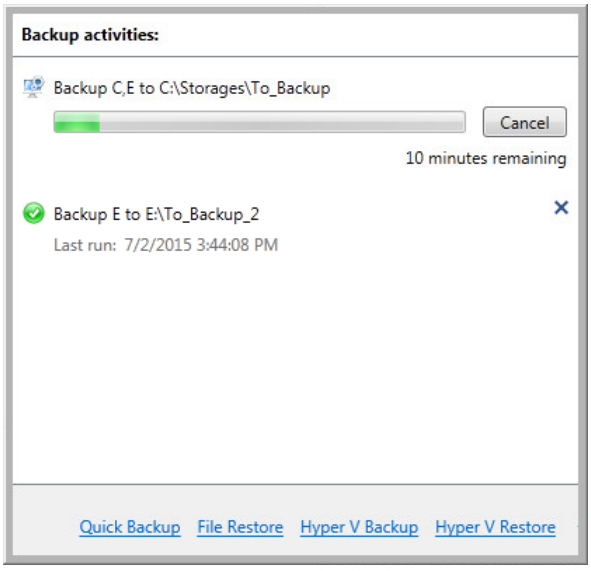

Backing up a Hyper-V guest machine

1. On the target machine please right click the Paragon's system tray application icon in the system tray, then select **Show activities**. In the opened Activities pane select the **Hyper-V Backup** hyperlink.

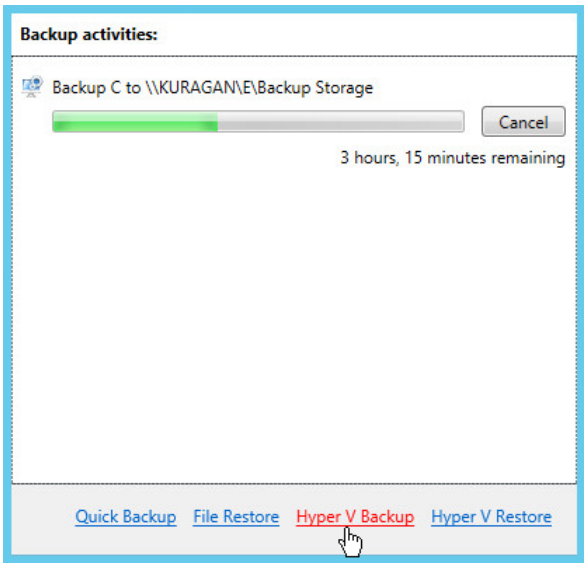

2. The wizard will list all virtual machines resided on the local Hyper-V Server. Click on the required machine (online or offline), then pick a timestamp you'd like to back up (if several available).

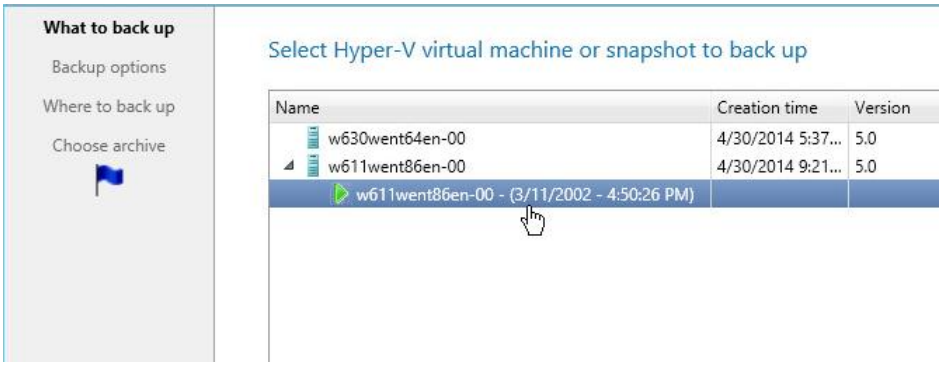

3. By default, the wizard is configured to back up the specified virtual machine to a pVHD virtual container, which you can change to VHD or VHDX. If using VHD/VHDX as target backup format, you can attach the container to an existing Hyper-V virtual machine and OS will be launched successfully.

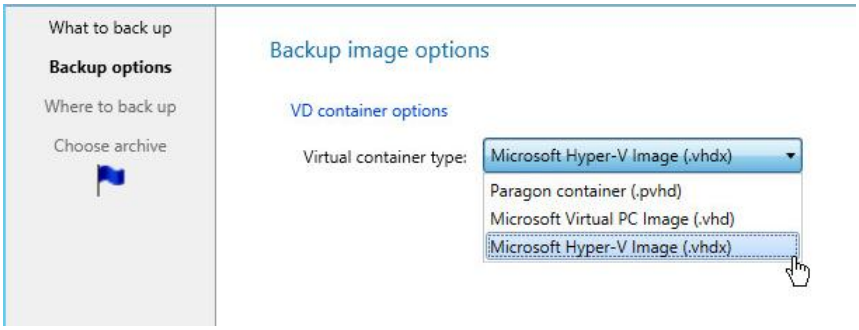

4. The wizard enables to save backup images either to a local folder, or a network share. Select the required backup location type in the corresponding menu, then specify a number of additional parameters. At this step you're also allowed to change the default backup name.

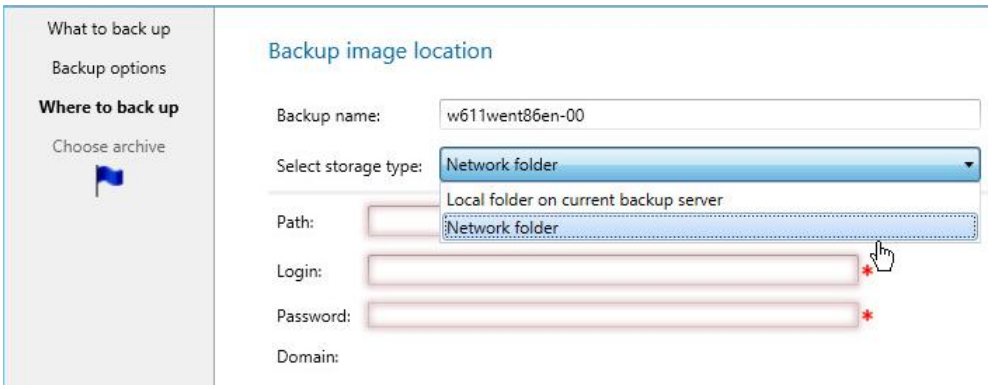

5. If it's not the first time you protect the target virtual machine to the specified location, you will be prompted to choose between creation of an incremental update or forcing a full backup.

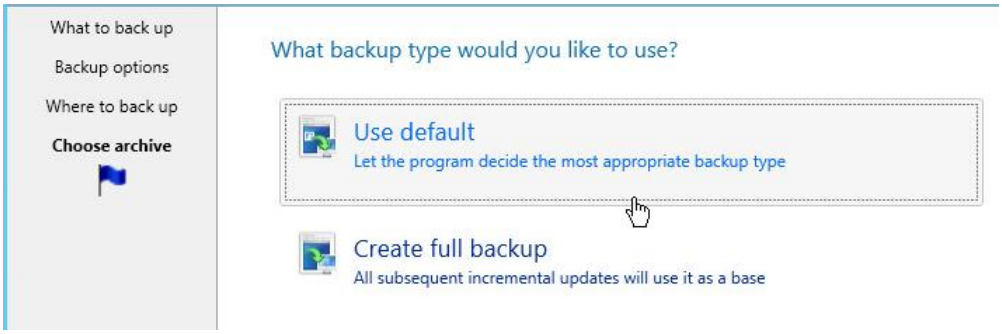

6. You will be informed on the operation start through a popup window. Monitor the backup progress through the Activities pane.

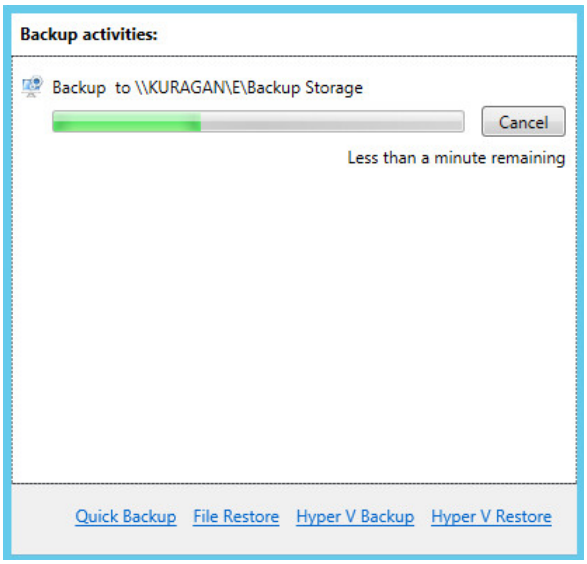

Restoring individual files from the host backup

1. On the target machine please right click the Paragon's system tray application icon in the system tray, then select **Show activities**. In the opened Activities pane select the **File Restore** hyperlink.

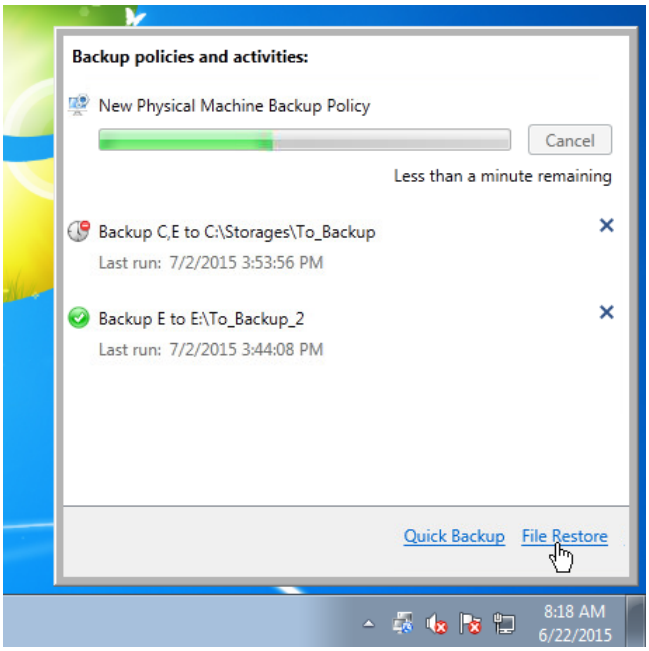

2. Specify a local folder or network share where backup images of the target machine are stored. Click here for more information.

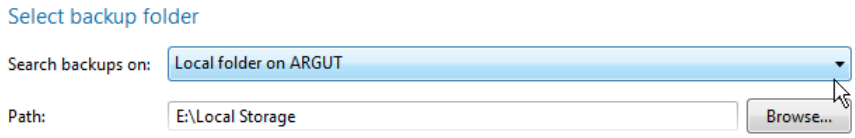

3. Choose a desired restore point. If there are too many items, please use the search pane to find the required image by date.

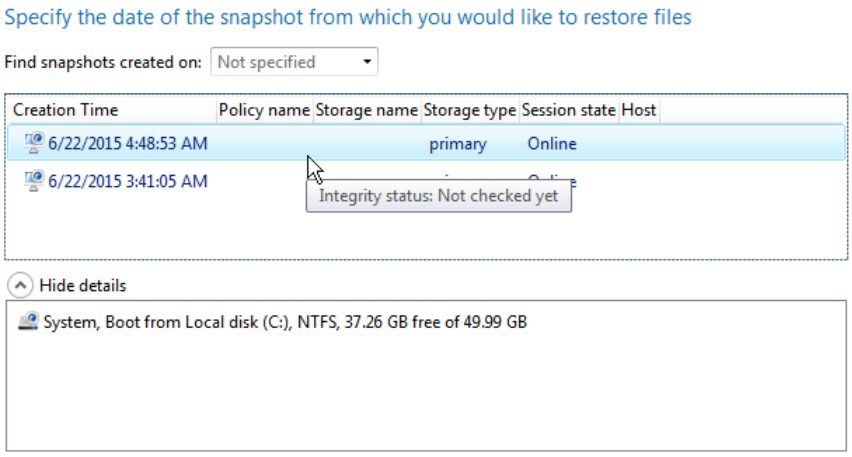

4. Find and mark files and/or folders, you'd like to restore.

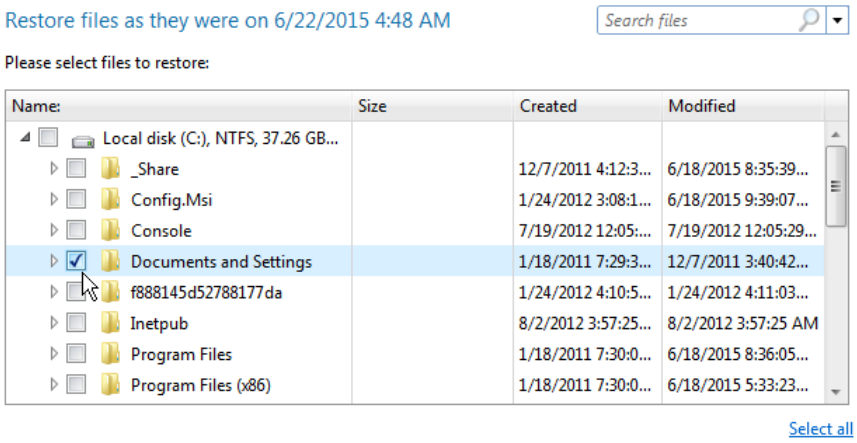

5. Specify where you'd like the selected backup data to be placed to (a local folder or a network share). If you'd like Windows Explorer to open in the specified folder once the operation is over, please additionally mark the corresponding option. Use the **Preserve directory structure** option to keep the original directory structure intact. Click **Start** when ready.

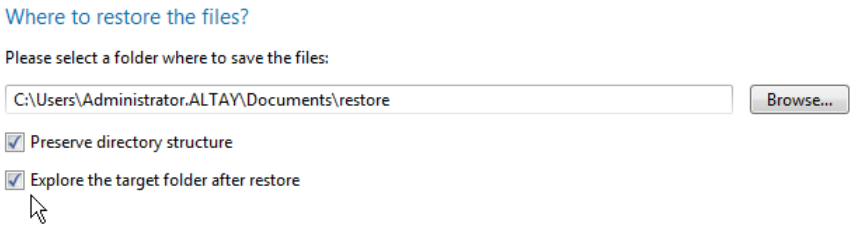

6. Monitor the operation progress. Click **Finish** when it's over.

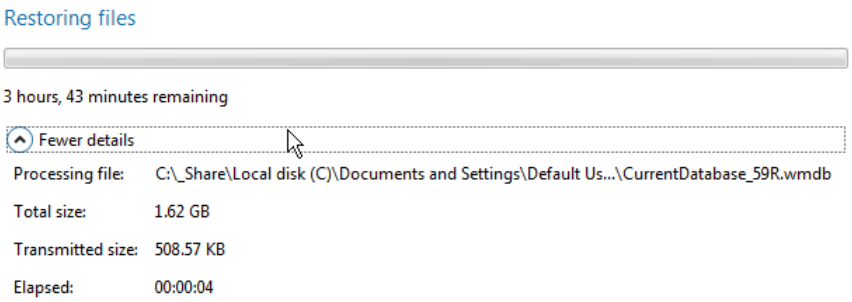

Restoring volumes from the host backup

Individual volumes can be restored with the help of the bootable WinPE media, which you can get from PPR Administrator. Please consult the corresponding scenario for more information.

Restoring a Hyper-V guest machine

1. On the target machine please right click the Paragon's system tray application icon in the system tray, then select **Show activities**. In the opened Activities pane select the **Hyper-V Restore** hyperlink.

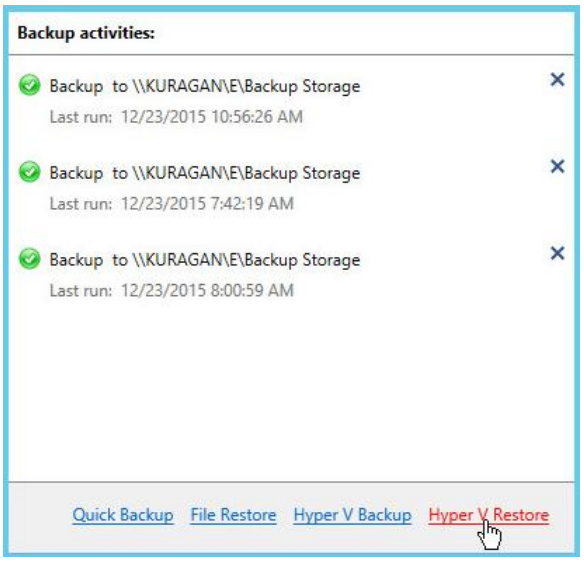

2. Specify a local folder or network share where backup images of Hyper-V guest machines are stored. Click here for more information.

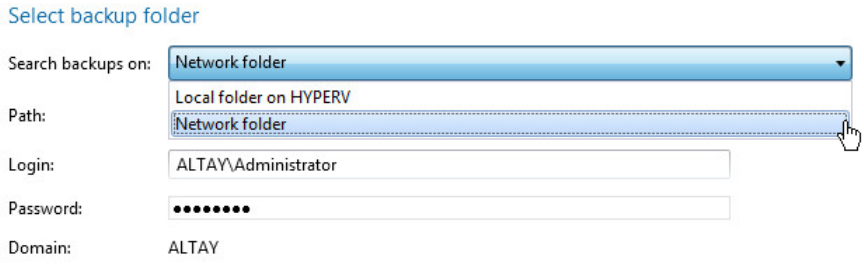

3. Select one of the previously backed up virtual machines from the list. If there are too many items, please use the search pane to find the required machine by name.

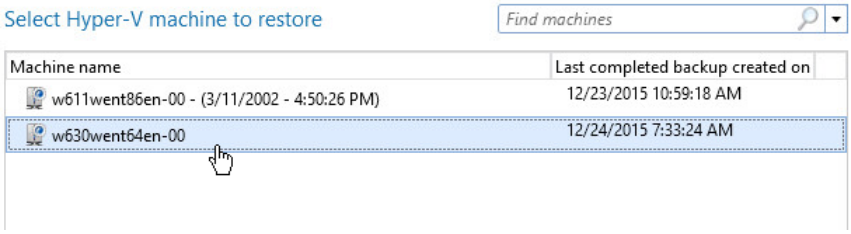

4. Choose a desired restore point. If there are too many items, please use the search pane to find the required image by date.

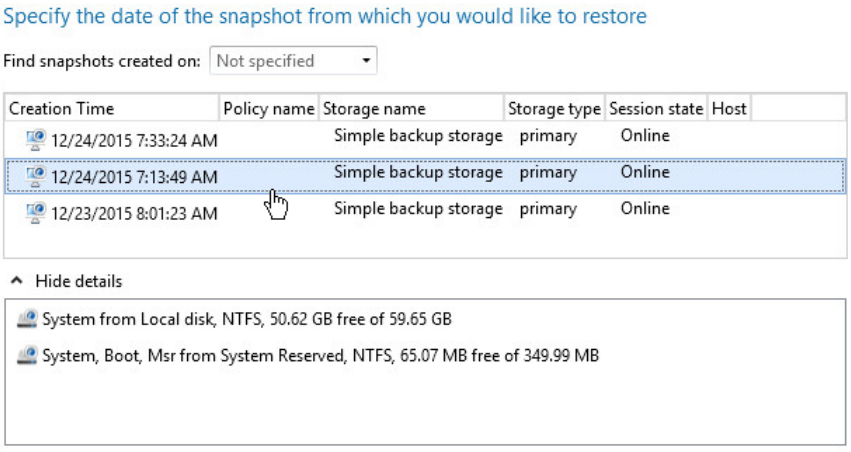

5. You've got the option to restore the specified backup data either to the original or a new location. Select the required mode to proceed.

6. If choosing the first option, just confirm the operation to start. Please note that all changes appeared on the machine after the specified time stamp will be irreversibly lost.

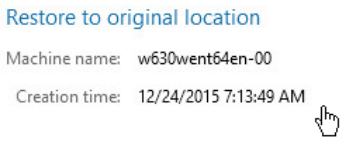

If attempting to restore to some other location, you will need to specify a local or network folder. By default, the postfix "_restored" will be added to the name of the virtual machine to differentiate it from the original one. However, you're free to give any name you like by entering it in the corresponding field.

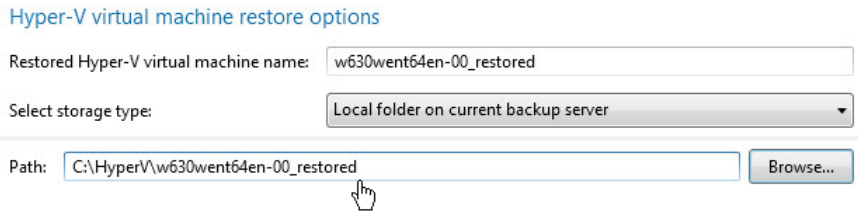

7. You will be informed on the operation start through a popup window.

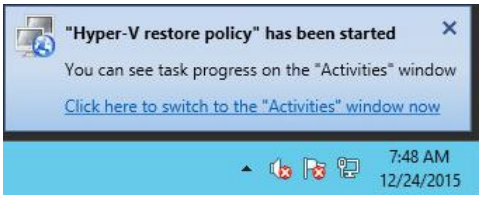

8. If restored to a new location you can see a new virtual machine in the Hyper-V Manager in the offline state once the operation is over. If restored to the original location, the original machine will be replaced by that from the backup image.

Extra Scenarios for WinPE

Adding specific drivers

To add drivers for specific hardware, please do the following:

1. Start up the target computer from the prepared WinPE media.
2. Select **Support & Tools**, then click **Add Drivers**.

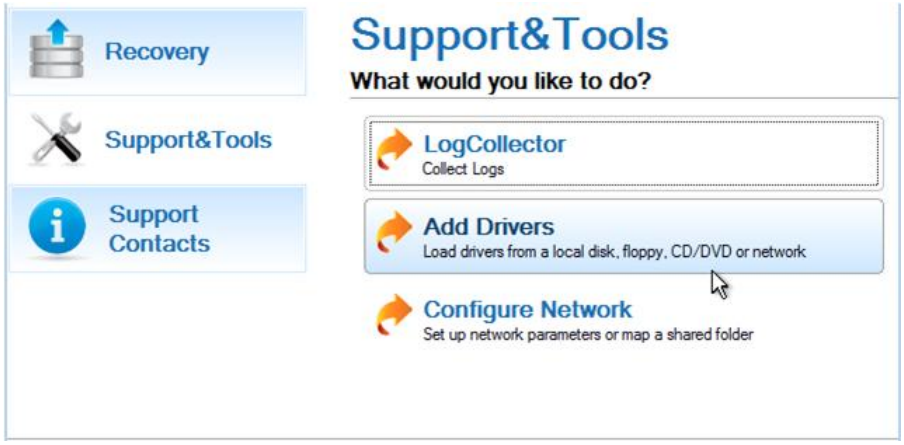

3. In the opened dialog browse for an .INF file of the required driver package located on a floppy disk, local disk, CD/DVD or a network share. Then click the **Open** button to initiate the operation

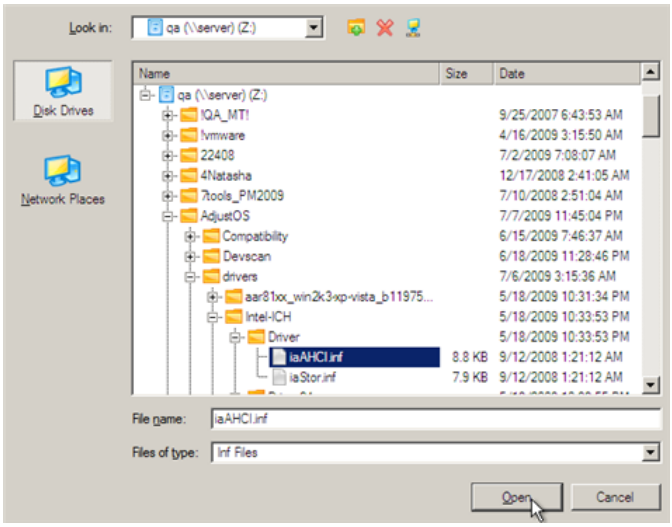

To know how to map a network share, please consult the Configuring network scenario.

4. You will be notified on the successful accomplishment of the operation. Click **Yes** to load another driver or **No** to close the dialog.

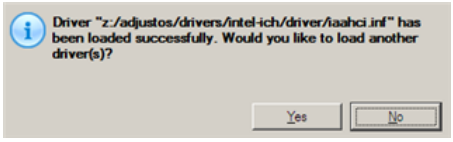

Configuring network

To manually set up a network connection and map a network share, please do the following:

1. Start up the target computer from the prepared WinPE media.
2. Select **Support & Tools**, then click **Configure Network**.

3. In the opened dialog provide an IP address, a network mask, default gateway, etc. for your network device.

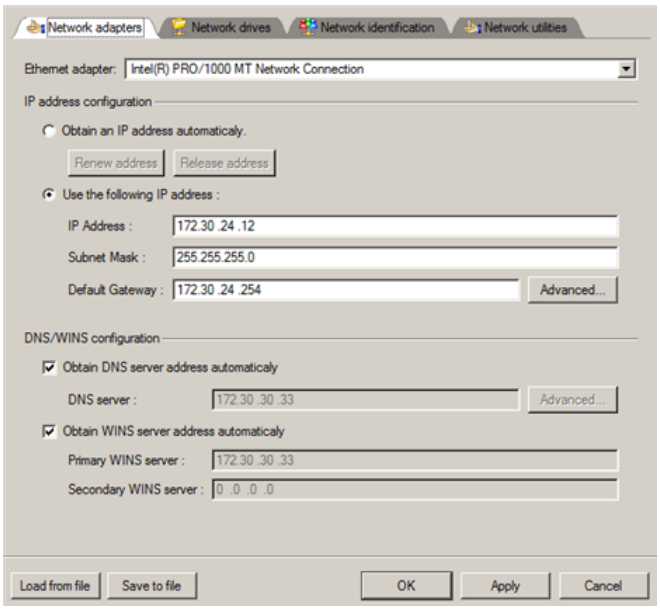

4. Click the **Network drives** tab to map a network share.

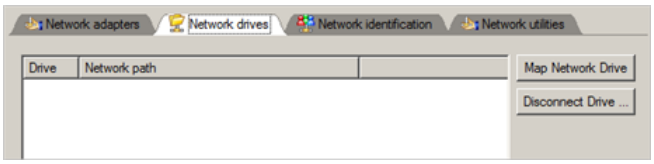

5. Click **Map Network Drive** and provide all the necessary information to map a network share in the opened dialog:

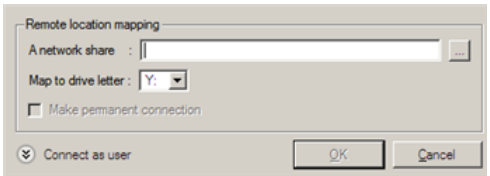

- Click the standard browse button [...] to browse for the required network share or manually enter a path to it;
- Define a letter from the pull-down list of available drive letters;
- Click the **Connect as user** button at the foot of the dialog page to specify a user name and password to access the selected network share if necessary.

By clicking **Disconnect Drive…** you can delete an existing network share if necessary.

6. Click the **Network identification** tab to change a network name of your computer (generated automatically) and a workgroup name.

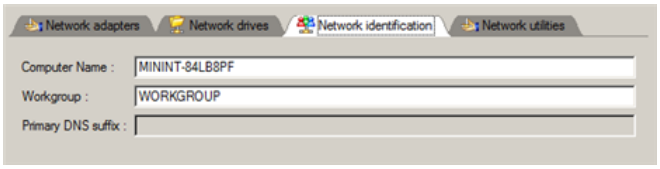

7. By default, the wizard saves all network settings in the netconf.ini file located on the WinPE RAM drive, thus it will only be available until you restart the computer. However, you can just once configure your network device and then save this file to some other destination, for instance a local drive, and this way avoid constant re-configuration, just by providing a path to it. So Click **Save to file** to save the netconfig.ini file to the required destination.

Network troubleshooter

Network Configurator includes a traceroute/ping utility that enables to get detailed information on particular routes and measure transit delays of packets across an Internet Protocol (IP) network. So with its help you can easily track down problematic nodes.

1. If you need to ping some network host, please select **Ping**, then type in the required IP address or its name. Click **Start** when ready.

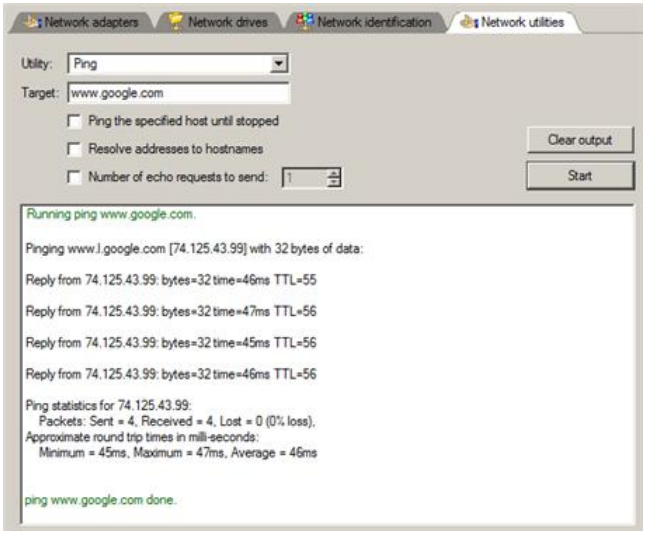

   - **Ping the specified host until stopped**. Mark the option to ping the chosen host for indefinite time;
   - **Resolve addresses to hostnames**. Mark the option to display hostnames instead of IP addresses.
   - **Number of echo requests to send**. By default the utility sends 4 echo requests, which you can modify however.

2. If you need to trace a route to some network host, please select **Trace route**, then type in the required IP address or its name. Click **Start** when ready.

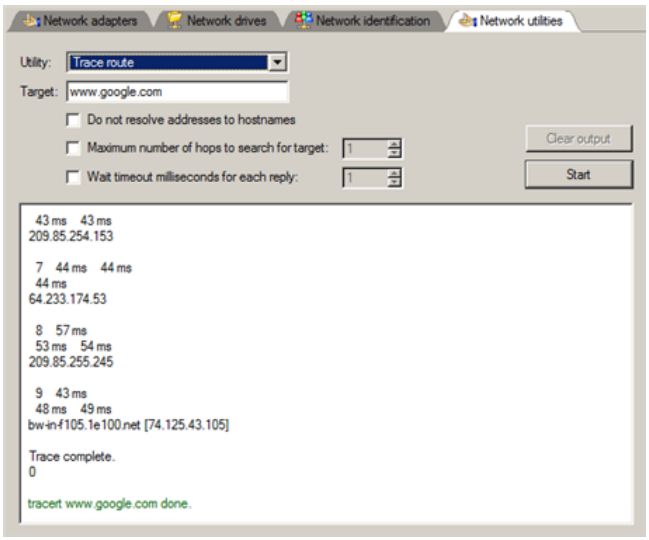

- **Do not resolve addresses to hostnames**. Mark the option to display IP addresses instead of hostnames.
- **Maximum number of hops to search for target**. By default the utility goes through maximum 30 hops when searching for the target host, which you can modify however.
- **Wait timeout milliseconds for each reply.** By default the utility waits 4 seconds for each echo reply message. If not received within the timeout, an asterisk (*) is displayed.

Protecting MS Exchange

Backing up Exchange Databases

The current version of the product can only back up all email databases of the specified Exchange Server.

Before you start please make sure all requirements for the target environment are met.

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **New physical machine backup policy**,

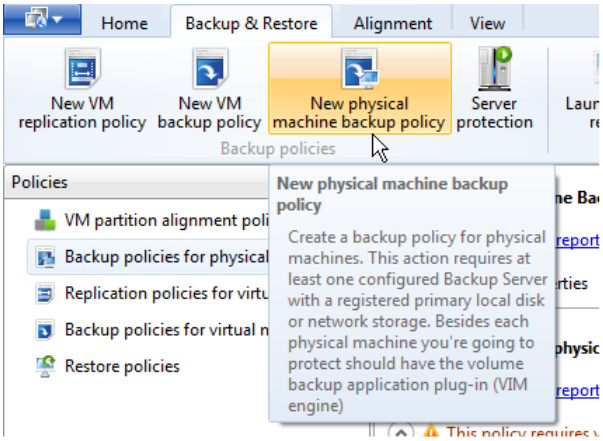

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **New VM backup policy**.

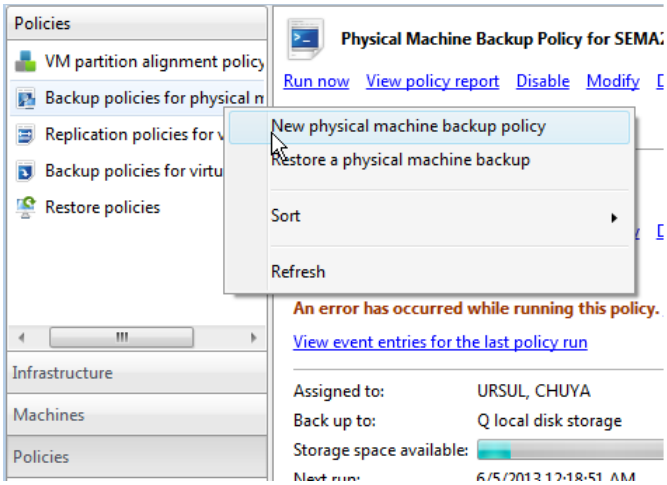

3. The opened dialog consists of four tabs that include a number of parameters:

   The first tab (Policy settings):

   - **Policy name**. Give it a catchy name.

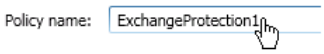

   - **Description**. Give a detailed description to the backup task (optional).

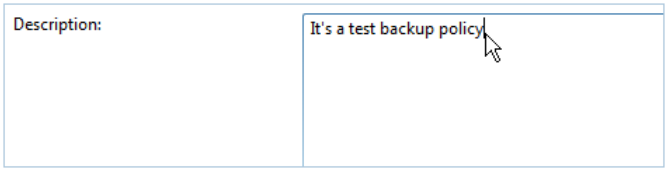

- **Back up to**. Select a backup server (if several), then the required primary storage from the popup list to place backup images to.

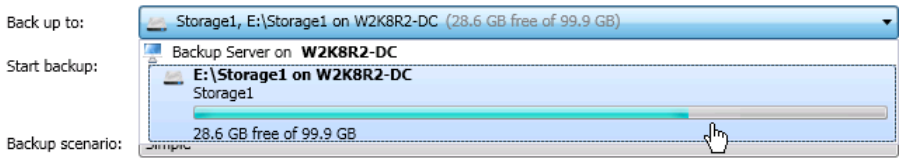

- **Start backup**. By default, no schedule is set for the backup policy. Click on the corresponding link to specify a timetable. Paragon's technicians recommend administrators to do a full backup once a day, while increments every 30 minutes – it enables to minimize restore time, backup storage requirements, and impact on Exchange.

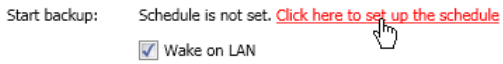

The opened dialog consists of two sections:

Basic scheduling

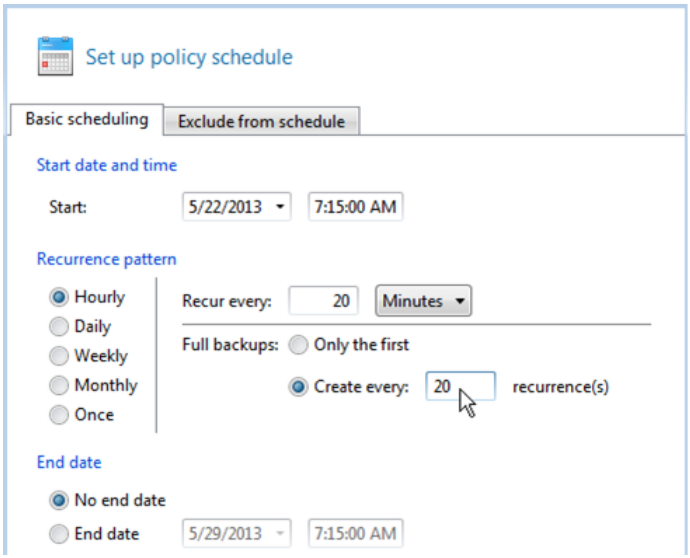

In this section you can set up a backup timetable. The minimal available update interval is one minute. By default, a full backup will be created once for every target machine, then only come incremental updates, which you can change however through the **Full backups** section.

Exclude from schedule

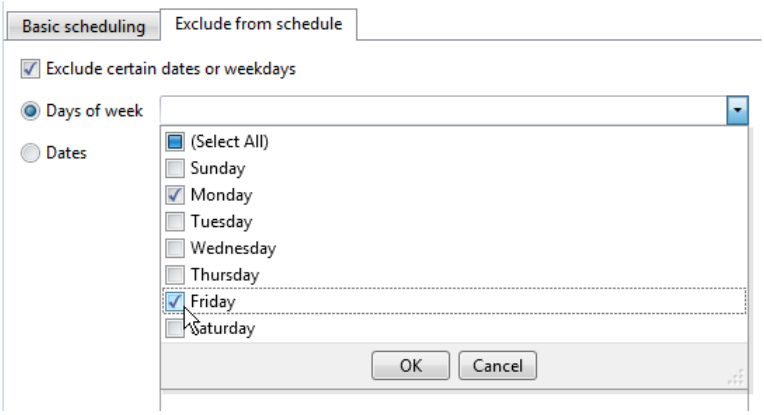

In this section you can specify days of week, or certain dates, when backup operations should not be accomplished.

If you schedule a backup task, please make sure time is the same on all computers. Anyway the operation will start according to time of Administration Server.

- **Wake on LAN**. Be default, the target physical machines will be automatically turned on to do backup through the Wake-on-LAN assistant.

- **Backup scenario**. In the current version of the product only one backup scenario is supported (simple).

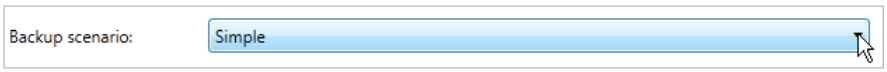

- **Data retention options**. Here you can specify a custom backup data retention mode that will be taken into account for the created policy only.

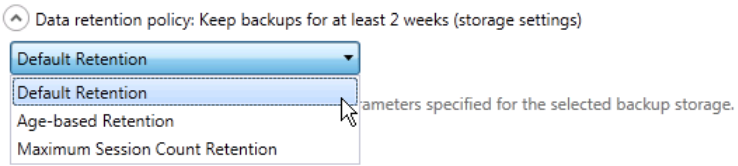

  - **Default Retention** to use data retention parameters specified for the selected backup storage.
  - **Age-based Retention** to define how often backups of target machines processed by this policy should be checked for exceeding the maximum amount of space they can take (10GB by default). On exceeding the set value, backup chains will be thinned out starting from the oldest backup images.

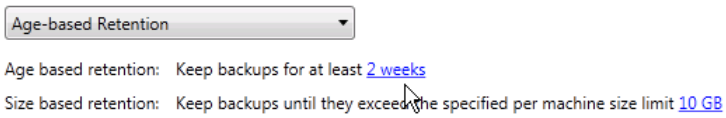

  - **Maximum Session Count Retention** to define the maximum number of backup sessions allowed for target machines processed by this policy. On exceeding the set value, backup chains will be thinned out starting from the oldest backup images.

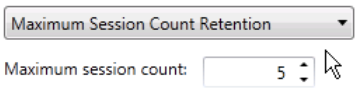

The second tab (Policy objects):

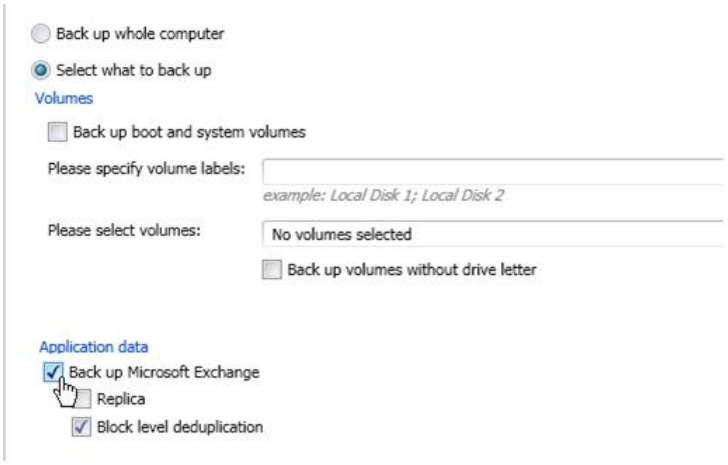

By default, entire machines will be protected. We don't need it, as in this case Exchange databases won't be treated at the application level. Click **Select what to back up** and then mark the option **Back up Microsoft Exchange**. If you'd like to protect Windows OS where Exchange Server is resided, please additionally specify the corresponding volumes manually, or just mark the option **Back up boot and system volumes**. Additionally you can specify whether to use the block-level de-duplication (highly recommended) and replica (it's only relevant for MS Exchange 2007, allowing use of database replicas during backup).

Specified in this section parameters will be applied to all target computers.

The third tab (Excludes). Here you can specify what data should be automatically ignored during backup. You can filter certain files or folders by creating masks. There are two types of filters:

- **Shared** that are applied to all physical backup policies. Click here for more information.
- **Private** that are created and applied to the current backup policy only.

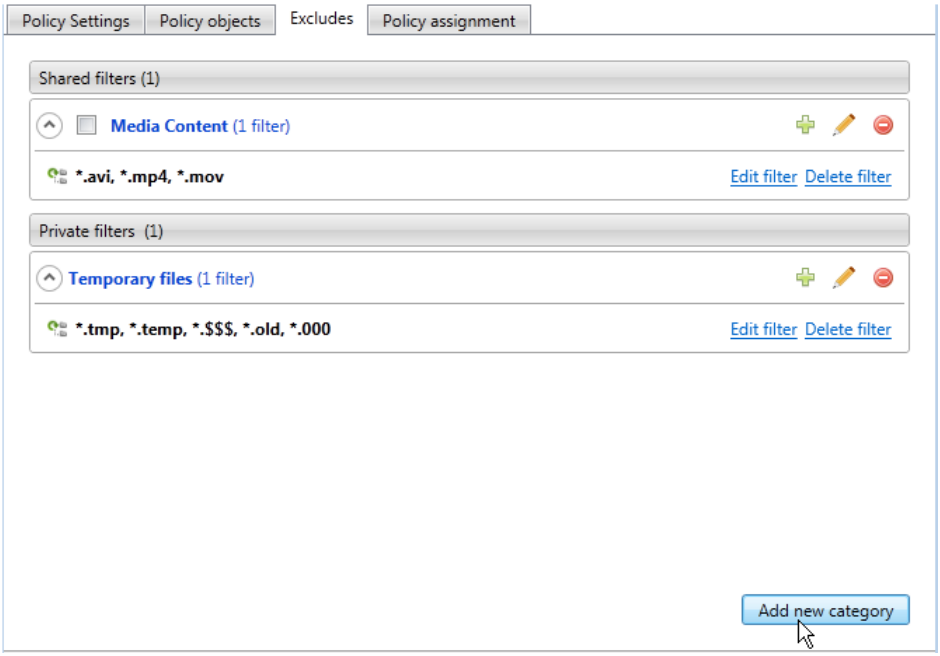

Exclude filters assigned at this stage will be applied to all target machines processed by this policy.

The fourth tab (Policy assignment). In this section you should specify a machine where MS Exchange is resided.

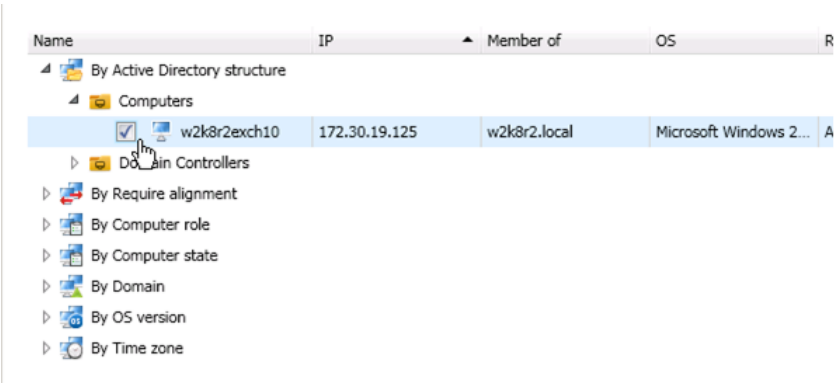

You can only specify a machine that already joins the infrastructure with the roles of 'Exchange Server 2007/2010 plug-in' and 'Volume backup application plug-in' (optional).

When you're ready with all parameters, click **Save & Validate** to complete creation of the backup policy. By default there will be used the fast level of validation, which you can change by clicking on the arrow button.

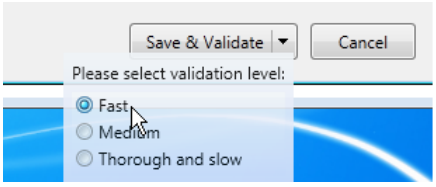

Let's see how three validation levels differ:

- **Fast**. It includes check up of all policy rules and their parameters and availability of the backup storage.

- **Medium**. It includes connection to the target machine to scan for specified backup objects as well as connection to the backup storage to retrieve metadata from it.
- **Thorough and slow**. It includes creation/deletion of snapshots of target virtual machines, creation of an uncompleted backup session and data items in the backup storage without opening data streams and data copying.

4. Validation of the backup task will be initiated immediately. You will be informed on the operation start through a popup window.
5. Additionally we recommend you to create a notification sub-policy for easier monitoring.
6. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information.
7. When the backup task is over, its status will be updated.

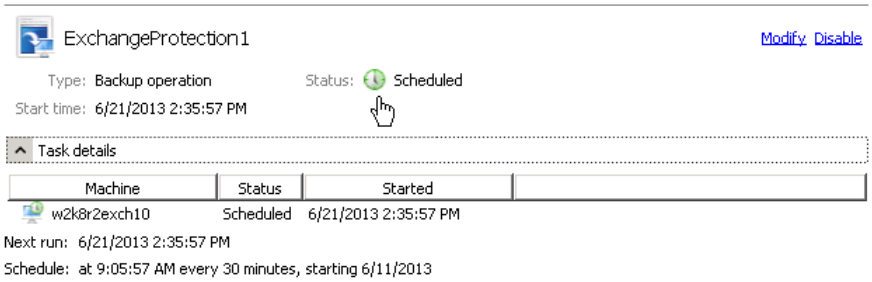

8. You can manage created backup policies through the **Policies** pane.

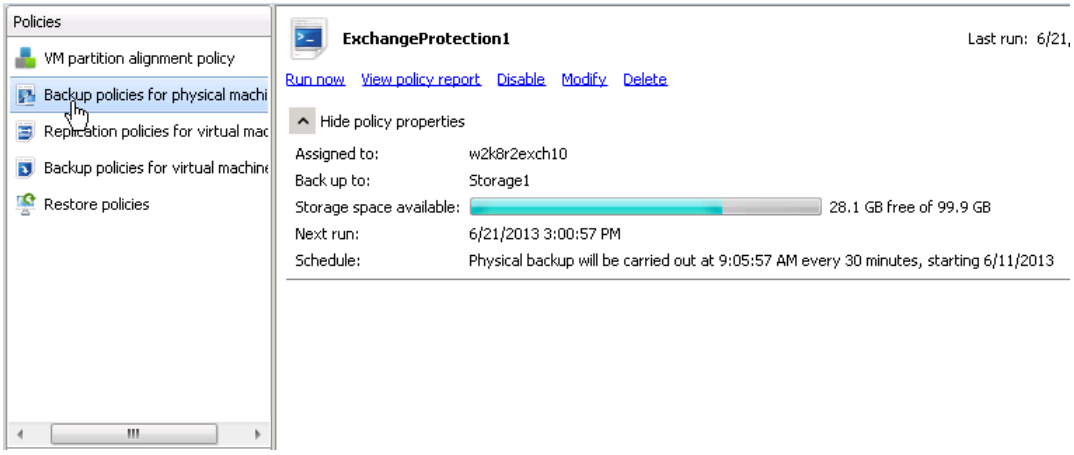

- **Run now**. Use this option to force launch of the required policy. If it's not the first time you commit a backup policy, you will be offered to choose the required backup mode.

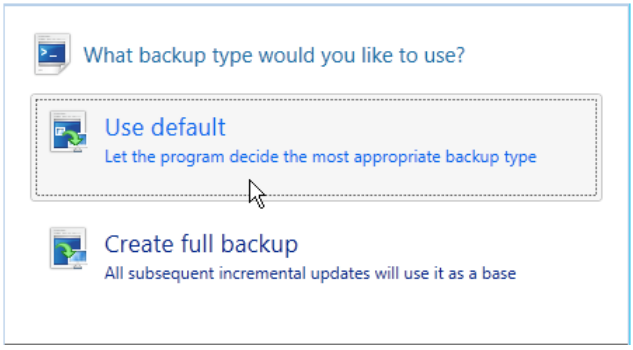

- **View policy report**. Use this option to get detailed information on all policy launches of the required policy. Here you can see when and with what result (succeeded or failed) each policy launch completed. Click **More…** next to a failed policy launch to see the reason.

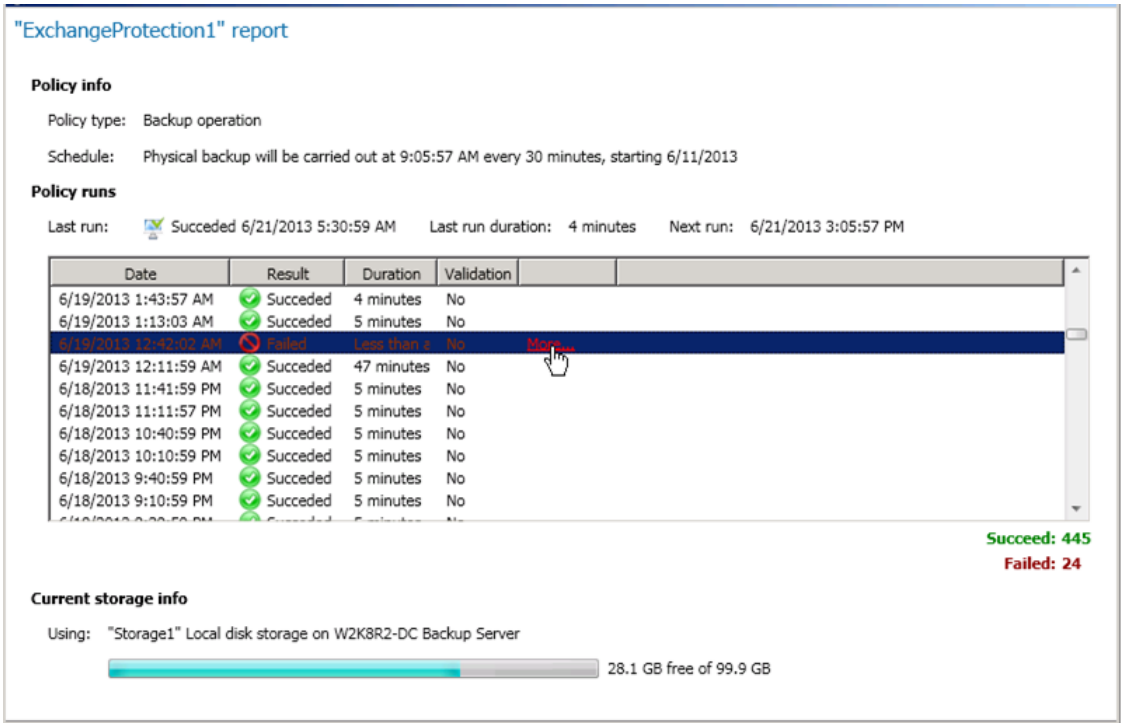

Select a failed object to see error details. If you need more information on the subject, please use the **Technical information** button to see infrastructure logs generated during the specified policy launch.

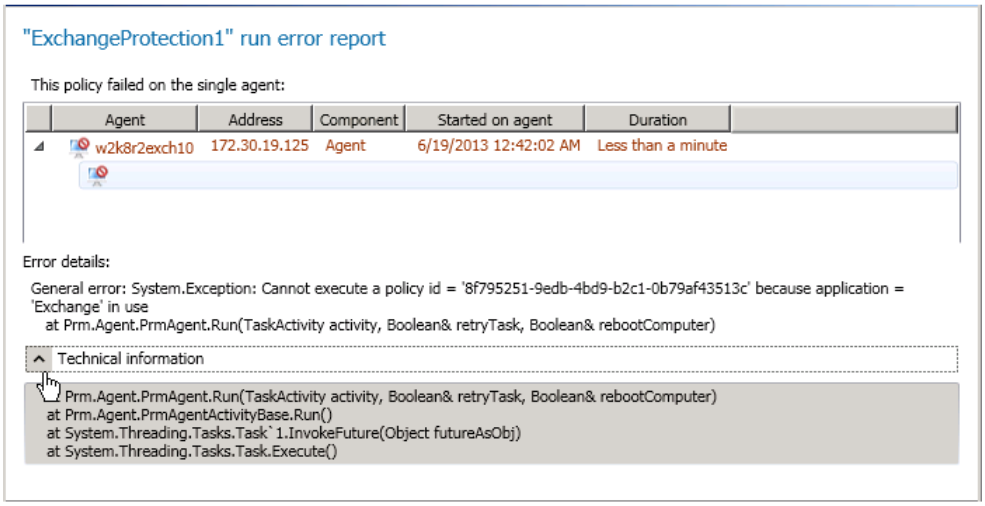

- **Disable**. Use this option to cancel the required policy.
- **Modify**. Use this option to change properties of the required policy. Please note the number of available properties will depend on the policy type (virtual backup, virtual replication, physical backup, etc.).
- **Delete**. Use this option to delete the required policy.

9. If you'd like to see and manage created backup images, use the **Browse storage** icon.

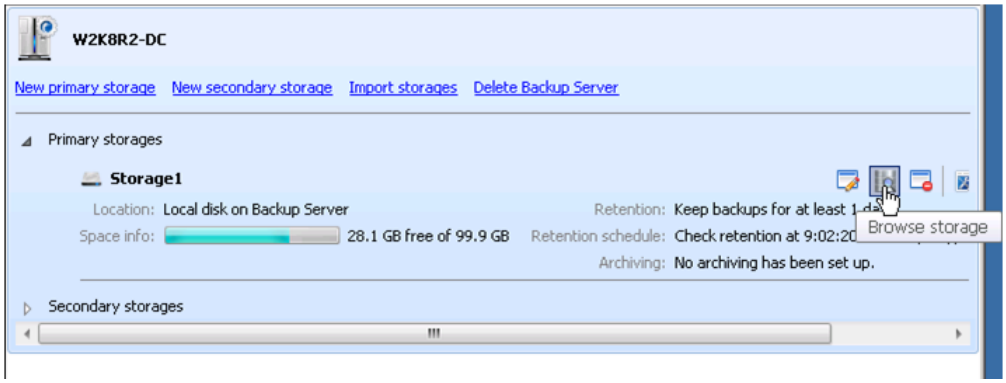

- In the opened dialog you can see a list of all protected machines with available backups. Use the **Find machines** field to quickly find the required object if necessary.

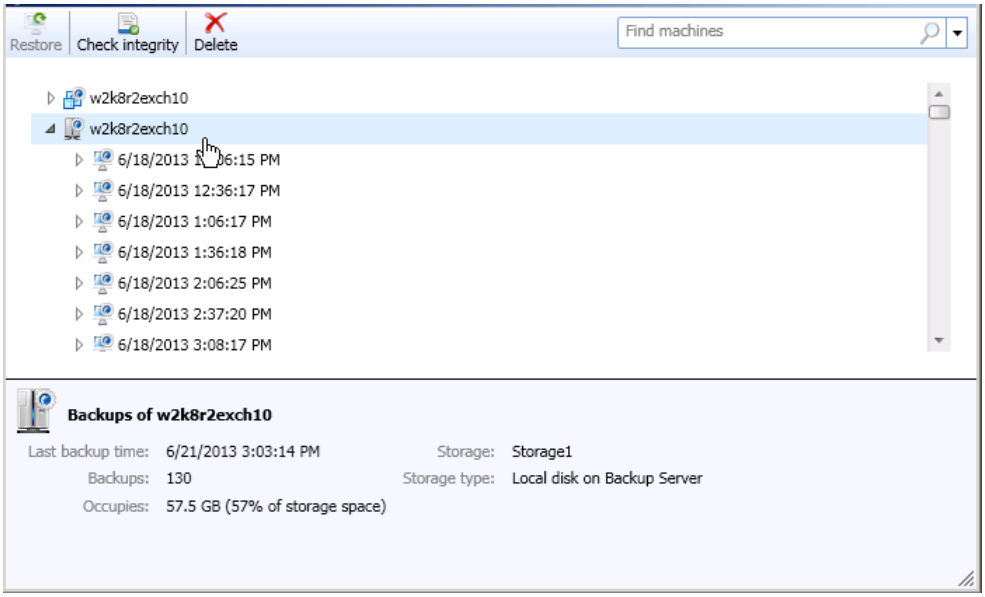

- Select the required machine, then one of available restore points to get details. Below you can see its type (full or incremental).

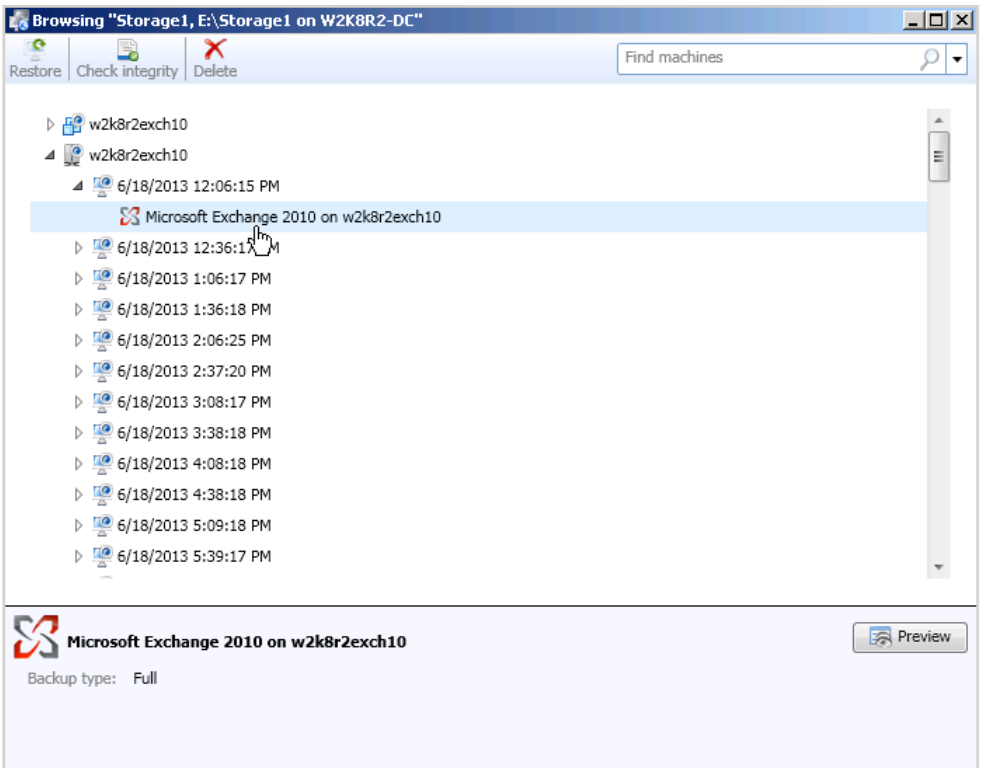

Click **Preview** to see its contents.

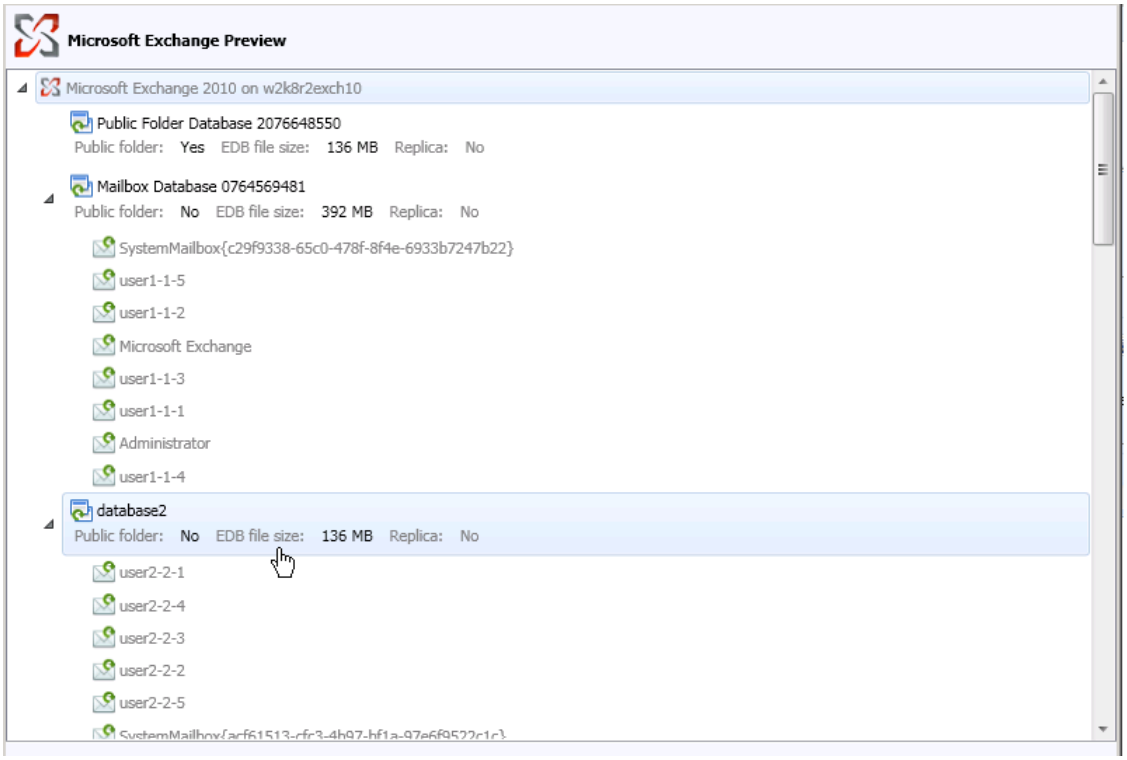

- Use icons above to:
  - **Initiate restore of the selected object**. To know more on the subject, please consult the corresponding restore scenarios.
  - **Verify the selected object for errors** by submitting a corresponding policy. If it turns out to be invalid, it will be marked by a special icon. We highly recommend you to check backup images for integrity before initiating restore.

- **Delete the selected object from the storage**. Please note if you delete an increment from somewhere in the middle of the incremental chain, the program will automatically initiate a data merging operation to keep the incremental chain valid. This action may take some time (depends on the amount of data of the deleted object) during which all members of the corresponding incremental chain will stay unavailable (grey out).

The same scenario can be accomplished through the System Protection Wizard.

Restoring Exchange Databases to Original Location

It's a basic restore scenario that enables to recover only selected or all backup Exchange databases to the original location. If restoring the latest backup image, you've got the option to additionally replay transaction logs for minimal possible dataloss.

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a physical machine backup**,

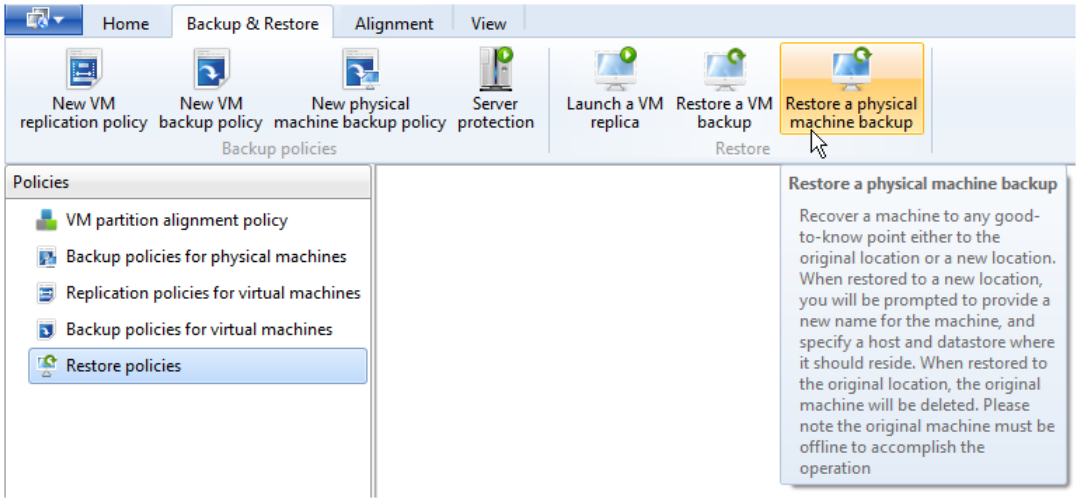

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a physical machine backup**.

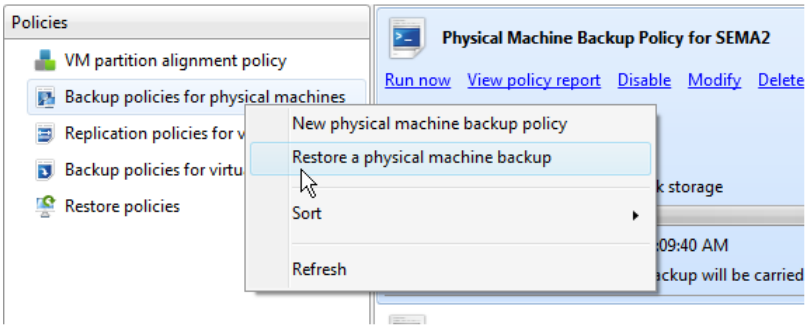

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name. Select a machine that hosts MS Exchange.

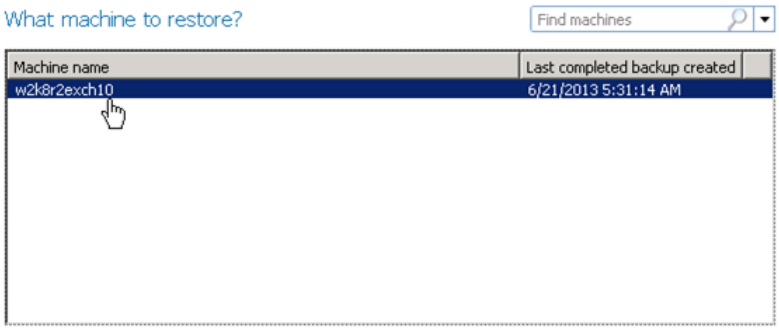

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

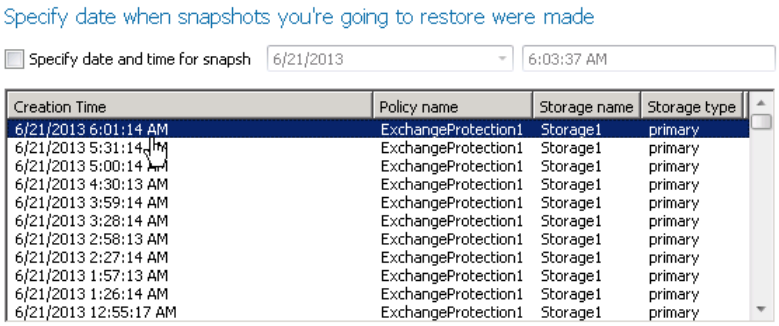

5. Since the selected restore point contains Exchange databases only, there's one available option: **Restore application data**.

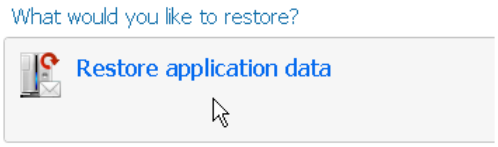

6. Select **Restore to original location**.

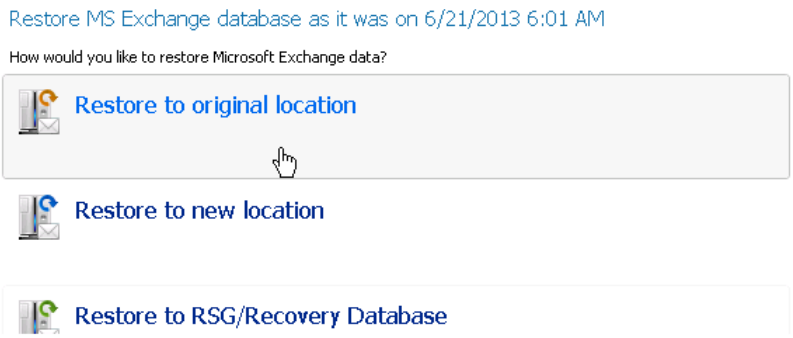

7. Select **Restore databases**.

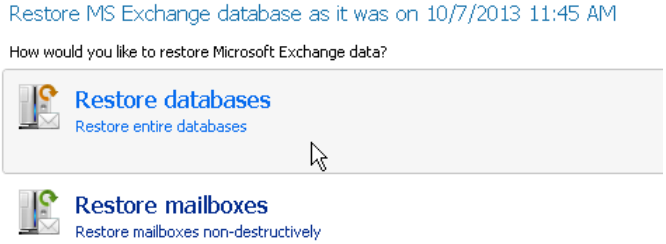

8. Mark the option **Restore all databases**, or specify the one you need to recover. If you have selected the latest increment to restore, there will appear an additional option **Use existing logs** (active by default), offering to replay transaction logs during the restore process. We highly recommend you to use this option, as it enables to reduce possible data loss to one email only. Click **Restore** to initiate the operation.

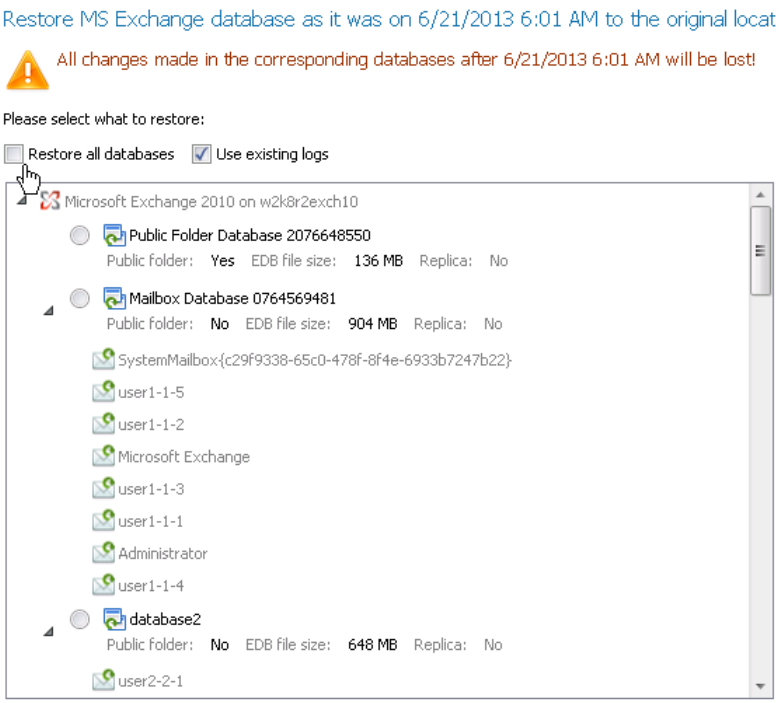

9. You will be informed on the operation start through a popup window.
10. Additionally we recommend you to create a notification sub-policy for easier monitoring.
11. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information.
12. When the restore task is over, its status will be updated.

Restoring Exchange Databases to New Location

Obviously restore of Exchange email databases may take hours, which is dangerous for any present day business. Inevitable downtime however can be practically avoided by using the so called dialtone databases during the process. The essence of this technology is that users of the problem Exchange Server keep sending or receiving emails, while their mailboxes are being restored. When the restore process is over all emails from the dialtone databases are moved to the restored databases, thus no information is lost. Microsoft provides instructions on how to use dialtone databases, but it's too complicated, involving many actions from the command-line PowerShell console. Paragon offers to do the same by going through a handy 5-step wizard, and even optimizes the process, getting unprecedentedly fast restore timings.

This scenario can only be used to restore a single Exchange database.

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a physical machine backup**,

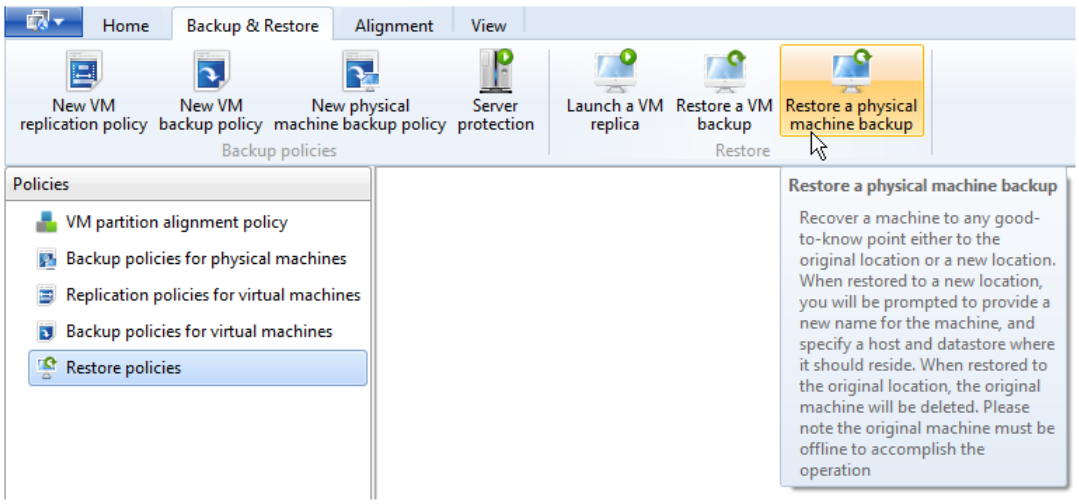

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a physical machine backup**.

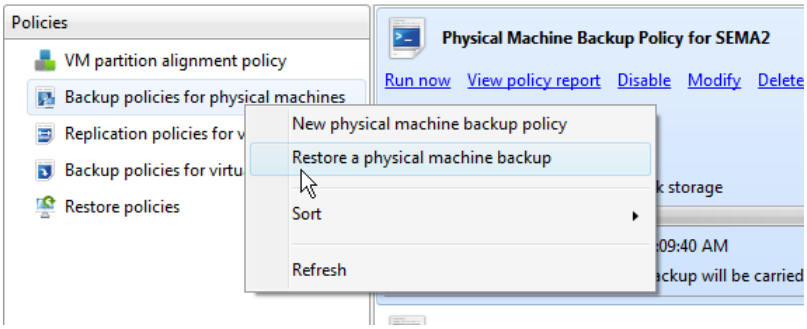

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name. Select a machine that hosts MS Exchange.

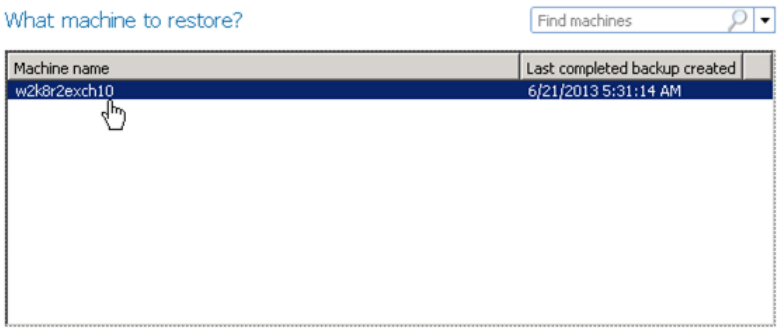

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

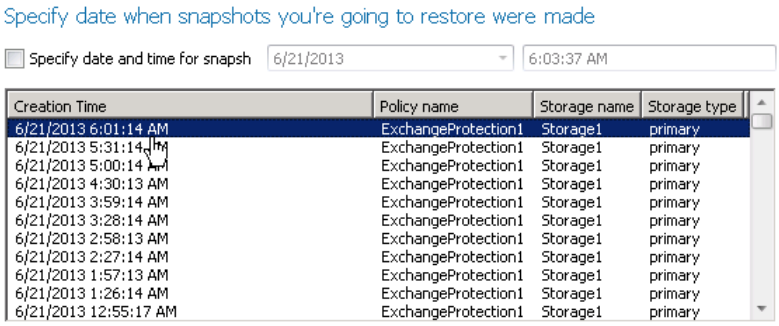

5. Since the selected restore point contains Exchange databases only, there's one available option: **Restore application data**.

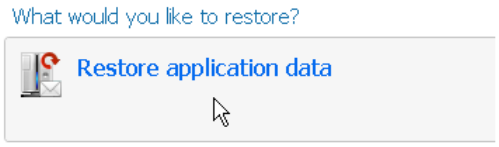

6. Select **Restore to new location**.

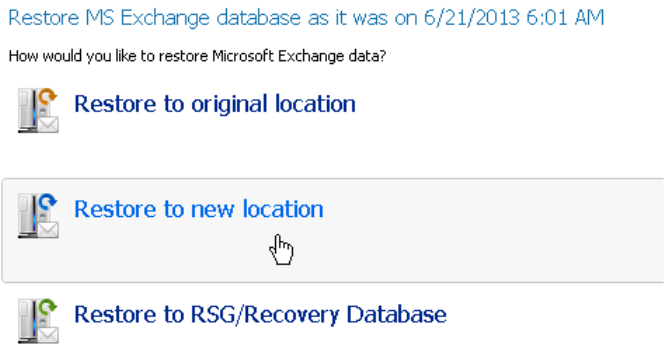

7. Select a machine where you'd like backup databases to restore. The wizard will only list machines that join the infrastructure and have the roles of **Exchange 2007/2010 application plug-in** (depends on the target Exchange).

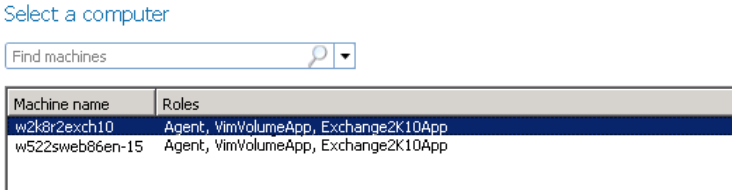

8. Select **Restore databases**.

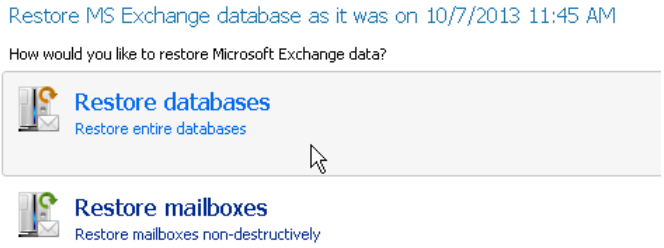

9. At this step you should specify a name for the restored database, its location and location of transaction logs.

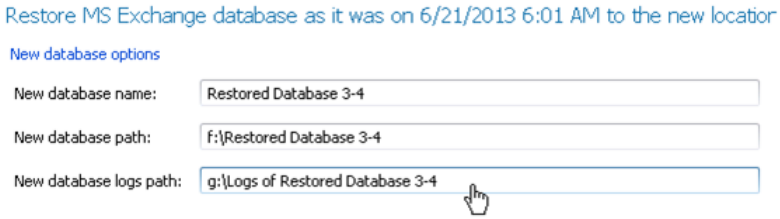

Besides here you can specify advanced restore options:

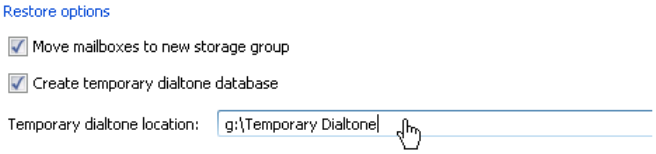

- **Move mailboxes to new storage group**. Mark the option if you'd like all mailboxes of the restored email database to move to a new storage group.
- **Create temporary dialtone database**. Mark the option if you'd like users of the failed storage group to keep sending / receiving emails while their mailboxes are being restored (highly recommended). When the restore process is over, all emails from the dialtone database will be automatically moved to the restored database and the dialtone database will be deleted.

In the current version of PPR there's no standard browse buttons, thus you should enter full paths manually.

10. Specify a database you need to recover. Click **Restore** to initiate the operation.

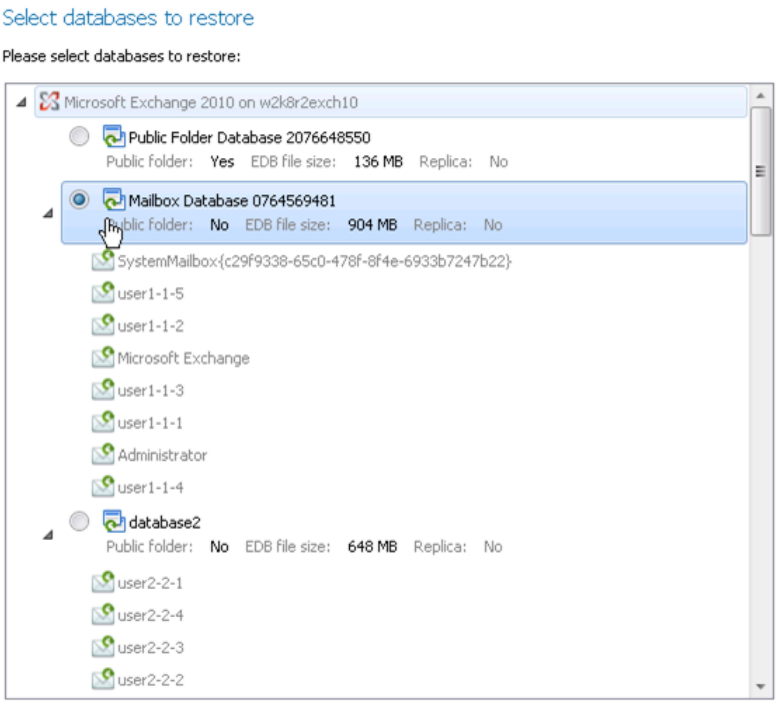

11. You will be informed on the operation start through a popup window.
12. Additionally we recommend you to create a notification sub-policy for easier monitoring.
13. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information.
14. When the restore task is over, its status will be updated.

Restoring Exchange Databases to RSG / Recovery Database

This is an advanced recovery scenario recommended by Microsoft. Unlike the previous one, it offers more options, but restore will take more time, as in this case it's not new data from the dialtone database that is moved to the recovery database, but vice versa. For more information please consult the Restoring an Exchange database to RSG/RDB chapter.

Restoring Mailboxes

This scenario has little to do with disaster recovery. Let's consider quite a typical situation when an Exchange administrator needs to restore a single mailbox – one of the employees left the company several months ago. According to the internal security policy, his email account was blocked and then deleted. Suddenly his former boss is requesting

to restore that mailbox as it contains very important information. Obviously, if the administrator doesn't have the option to restore at the mailbox level, he's in trouble. The same goes when important emails are deleted by accident.

PPR allows non-destructive restore of certain mailboxes. Mailbox contents can either be restored to the original location, provided none of the already existed email items are lost, or to a new location.

To restore one or several mailboxes from an Exchange database backup to a new location, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a physical machine backup**,

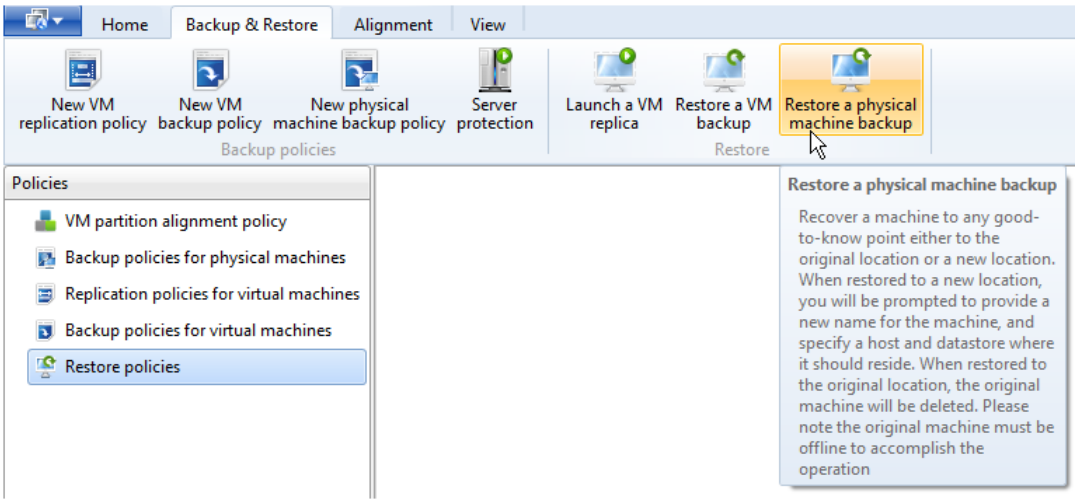

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a physical machine backup**.

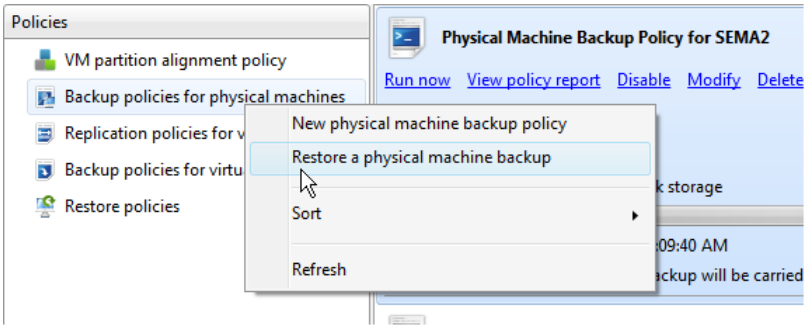

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name. Select a machine that hosts MS Exchange.

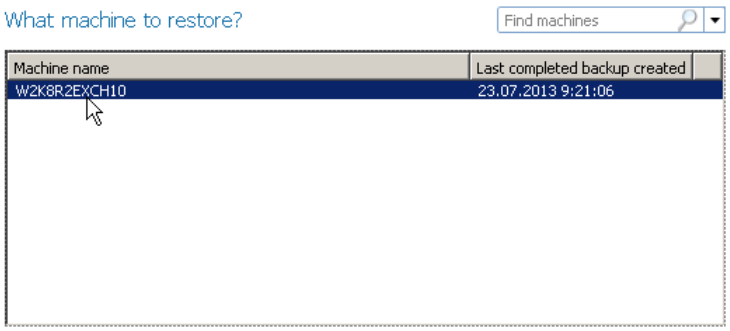

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

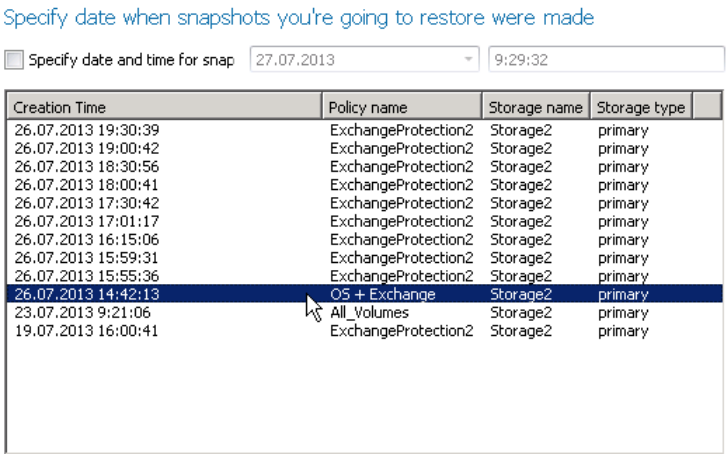

5. Since the selected restore point contains both, Exchange databases and system volumes, there are several available options. Please click on **Restore application data**.

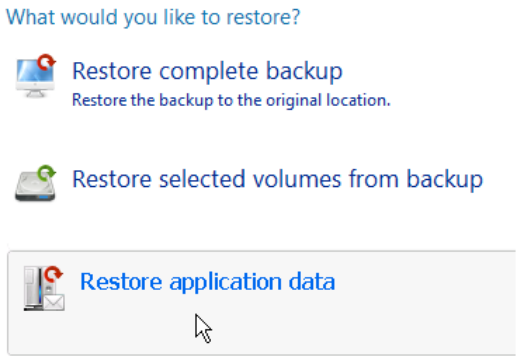

6. Select **Restore to new location**.

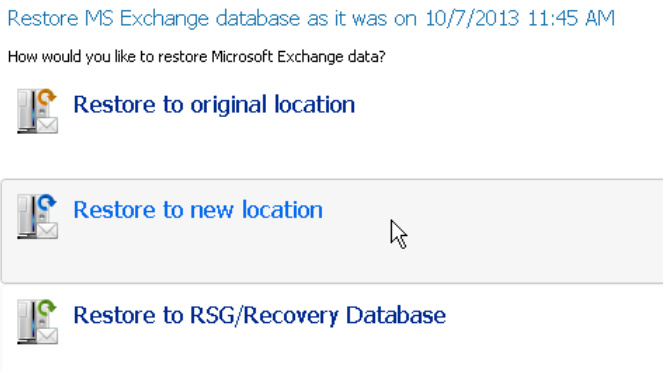

Restore of certain mailboxes can be done either to the original or a new location.

7. Select a machine where you'd like backup mailboxes to restore. The wizard will only list machines that join the infrastructure and have the roles of **Exchange 2007/2010 application plug-in** (depends on the target Exchange).

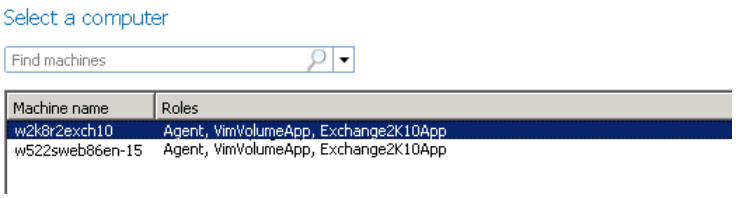

8. Select **Restore mailboxes**.

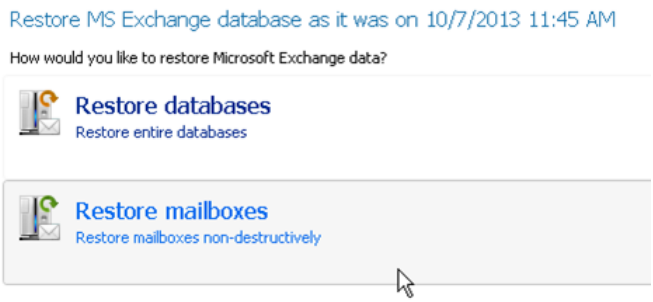

9. Specify a database that contains the required mailboxes.

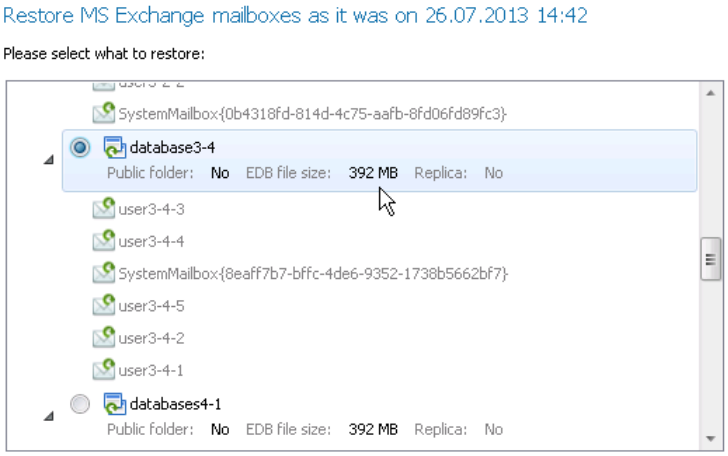

10. Specify one or several mailboxes you'd like to restore by marking checkboxes next to their names, or use the **Select all** option to select all mailboxes of the specified database in one action. To restore all selected mailboxes to a particular mailbox, just type in its name in the corresponding field or use the **Browse** button to navigate through all databases and their mailboxes of the specified Exchange Server. In this case there will be created a subfolder for each mailbox, named after them. However, you've got the option to specify your own name of the folder in the **Target Folder** field. In our example we restore one mailbox (user3-4-4) to another (user1-1-1). Click **Restore** to initiate the operation.

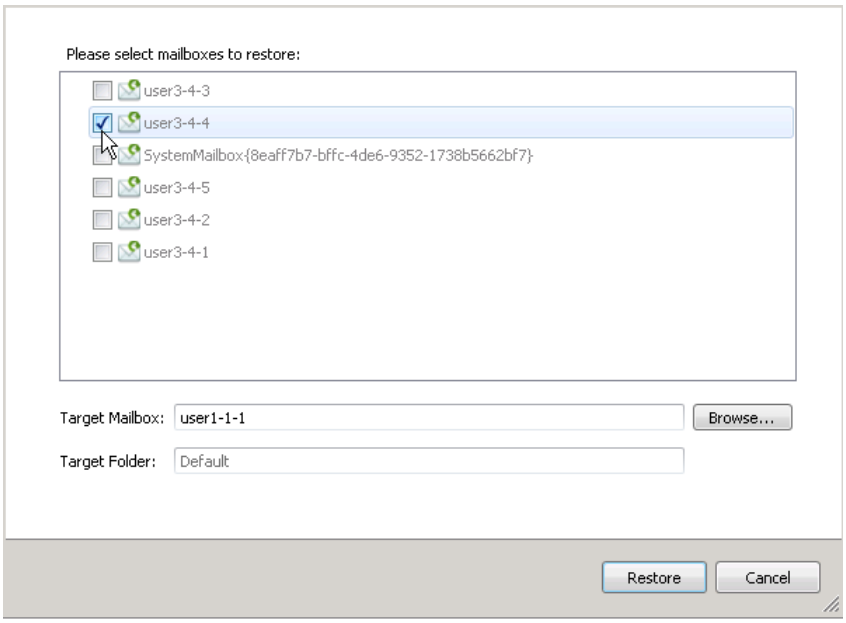

Browsing Exchange:

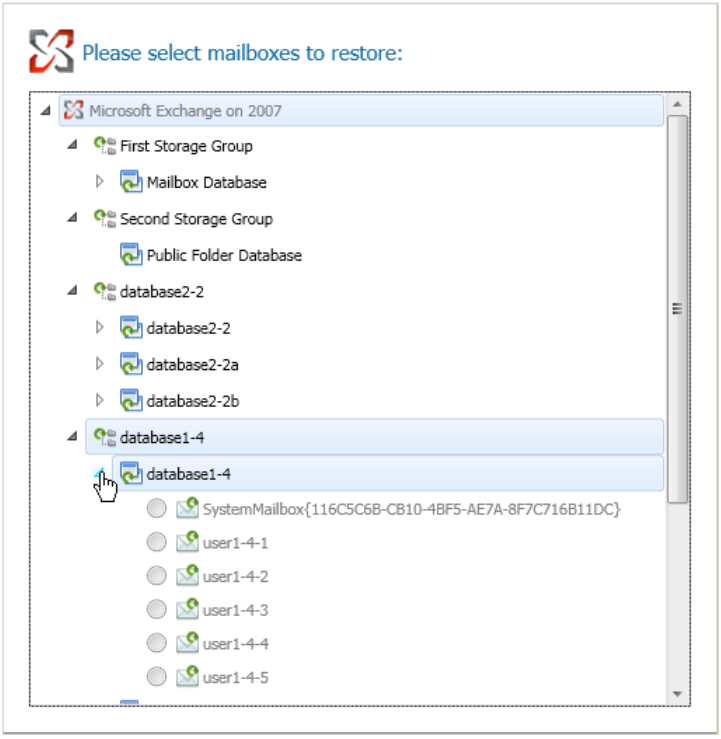

11. You will be informed on the operation start through a popup window.

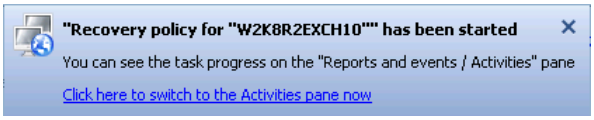

12. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information.

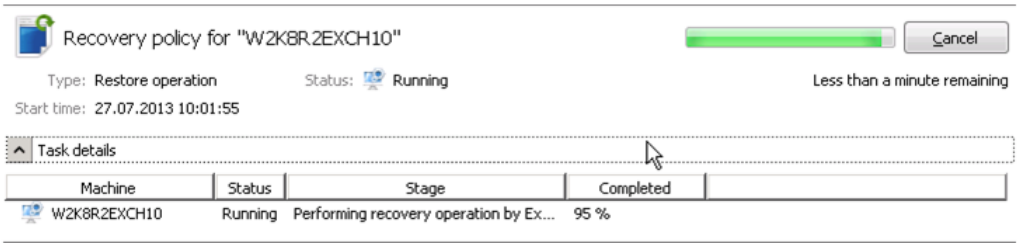

13. When the restore task is over, its status will be updated.
14. Now we open MS Outlook to make sure that mailbox **user3-4-4** was restored to mailbox **user1-1-1**, just exactly as specified.

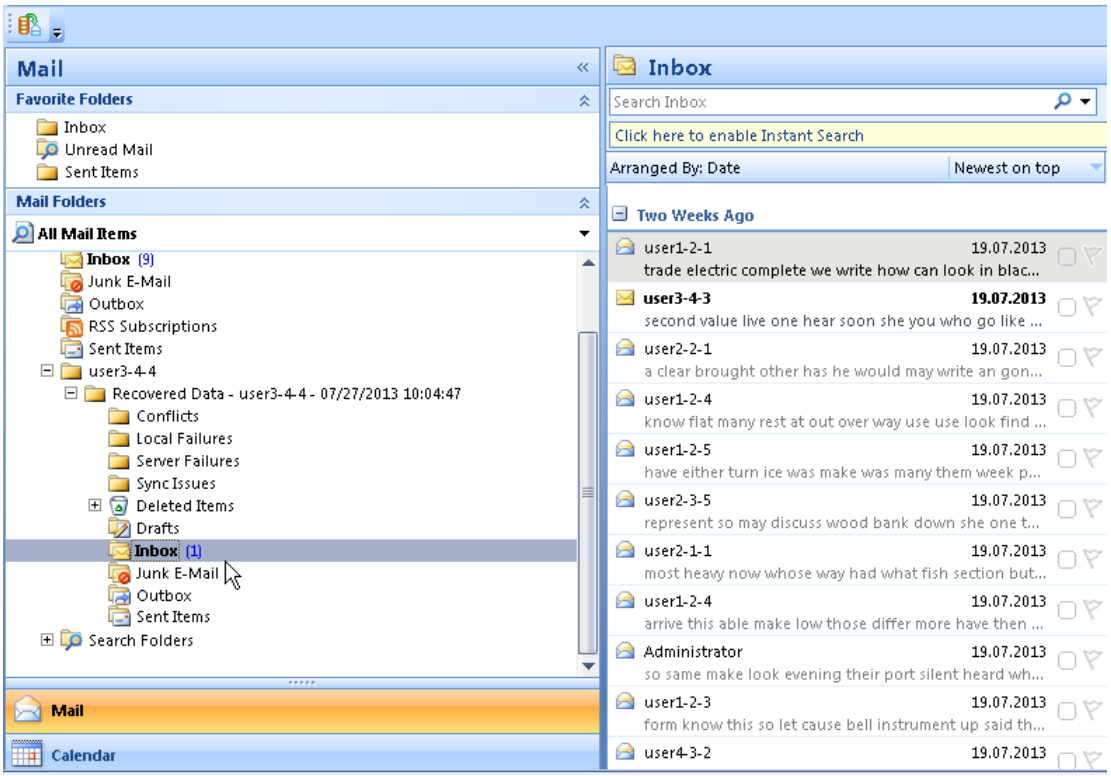

Using PPR and PEGR to Restore Emails from Exchange Databases

Now we will show you how PPR can be complemented with Paragon Exchange Granular Recovery to provide a rounded scenario of email items recovery from backup Exchange databases. This scenario includes several phases:

- Restoring an Exchange database to Recovery Storage Group (RSG) or Recovery Database (RDB);
- Preparing Exchange RSG/RDB for connection through PEGR;
- Installation of PEGR;
- Restoring email items of Exchange RSG/RDB in MS Outlook with embedded PEGR.

Restoring an Exchange database to RSG/RDB

To be able to safely work with contents of an Exchange database through PEGR and MS Outlook, first you need to get its backup copy. The most efficient way to get it is to restore the required database to a new location or to RSG/RDB. We choose the second option:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Backup & Restore** ribbon, then select **Restore a physical machine backup**,

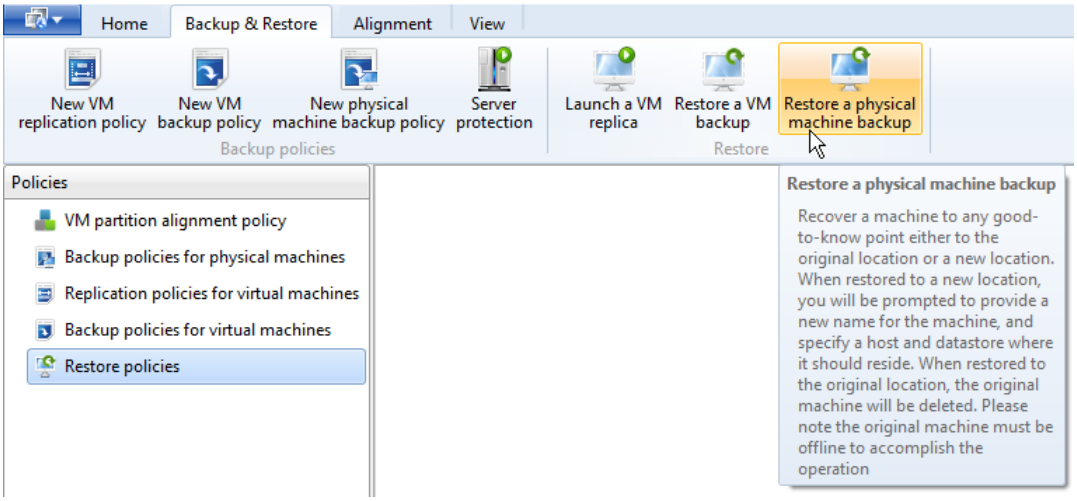

or go to **Policies >** right click on the **Backup policies for physical machines**, then select **Restore a physical machine backup**.

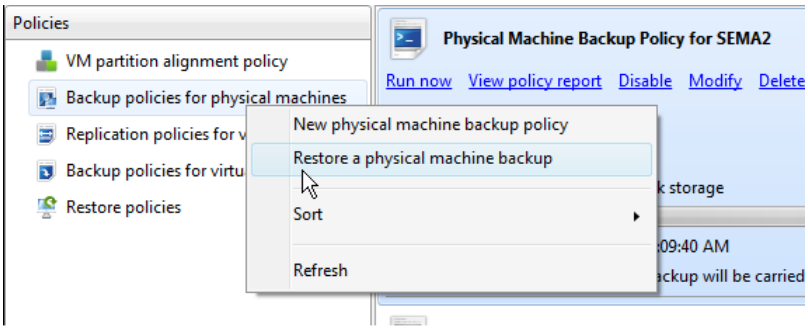

3. The opened wizard will first prompt you to select one of the backed up earlier physical machines. If there are too many items on the list, please use the search pane to find the required machine by name. Select a machine that hosts MS Exchange.

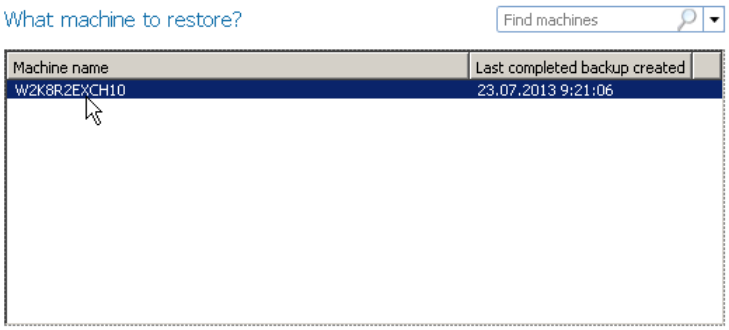

4. Then you will need to choose a desired restore point, if several. If there are too many items on the list, filter the list by marking the checkbox **Specify date and time…**, then providing the required date and time.

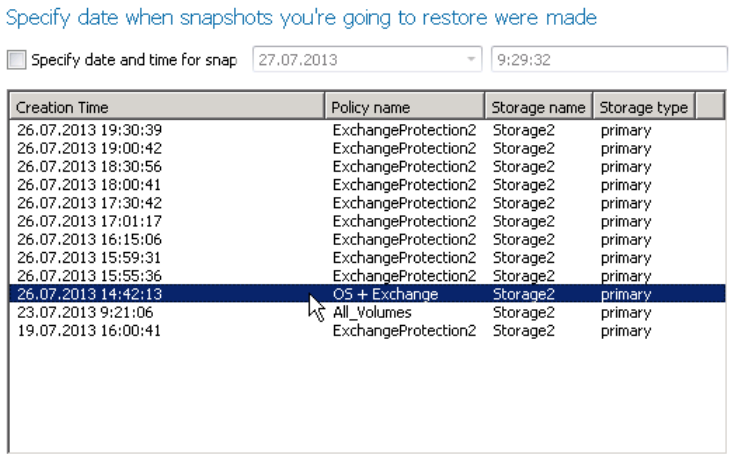

5. Since the selected restore point contains both, Exchange databases and system volumes, there are several available options. Please click on **Restore application data**.

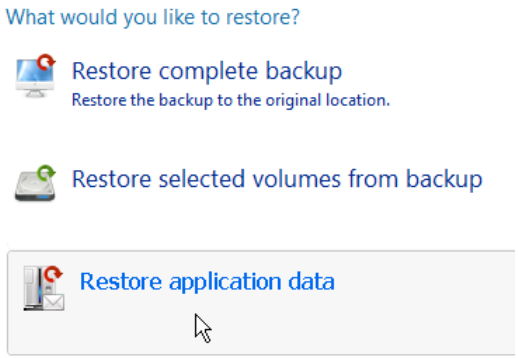

6. Select **Restore to RSG/Recovery Database**.

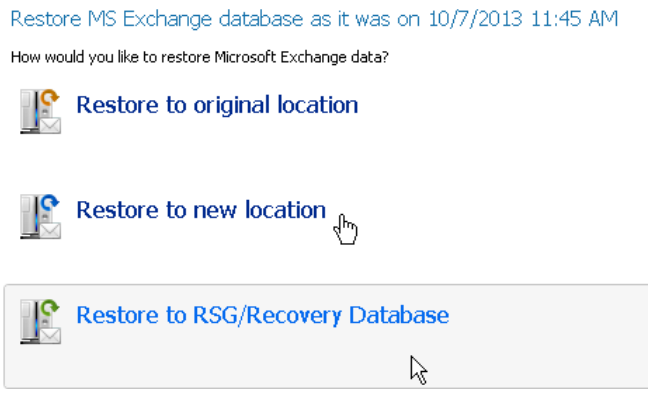

You can also use the 'Restore to new location' option to get the required database copy. To know more on the subject, please consult the Restoring Exchange Databases to New Location chapter.

7. Select a machine where you'd like backup databases to restore. The wizard will only list machines that join the infrastructure and have the roles of **Exchange 2007/2010 application plug-in** (depends on the target Exchange).

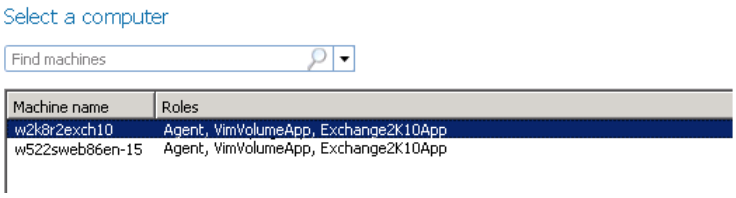

8. At this step you should specify location of the resulted RSG/RDB and transaction logs or mark the corresponding option to use the default locations. For MS Exchange 2010 it is **C:\Program Files\Microsoft\Exchange Server\V14\Mailbox\RDB…**

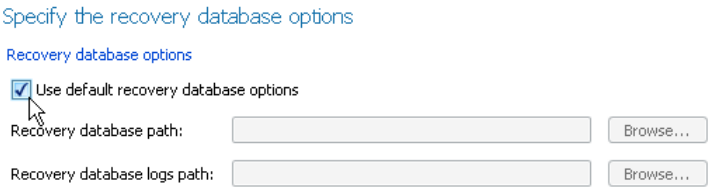

9. On the next page please use the default options.

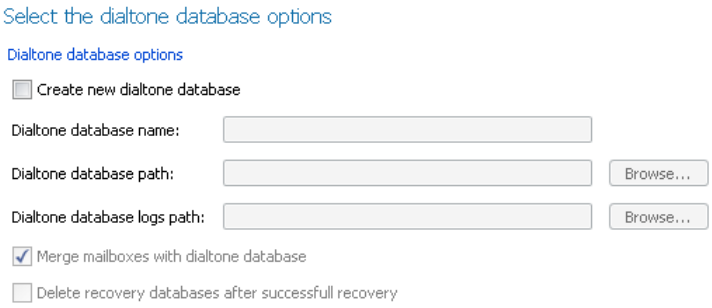

10. Specify a database you need to recover. Click **Restore** to initiate the operation.

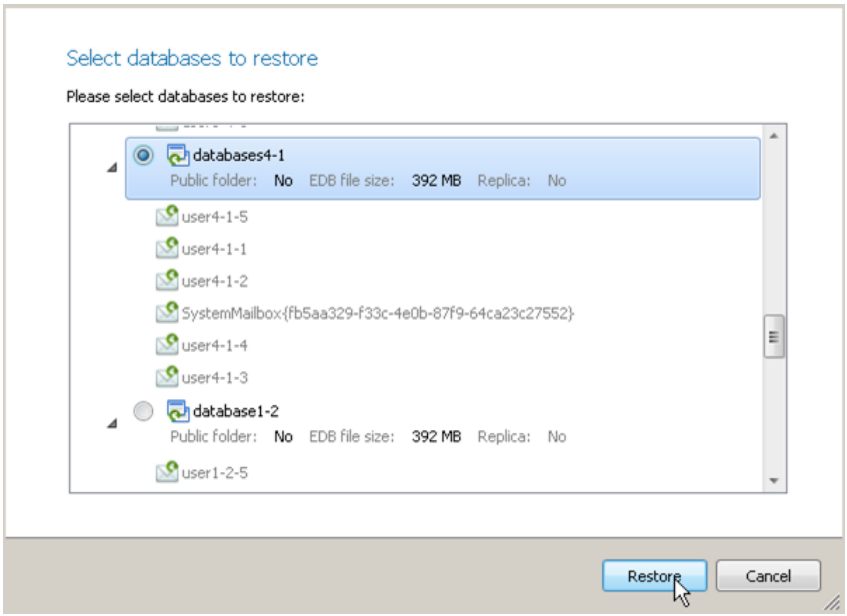

11. You will be informed on the operation start through a popup window.
12. Additionally we recommend you to create a notification sub-policy for easier monitoring.
13. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information.

14. When the restore task is over, its status will be updated.

Preparing Exchange RSG/RDB for connection through PEGR

Once you've got the required database as RSG/RDB, you should dismount it to be able to connect through PEGR:

1. Open Exchange Management Console

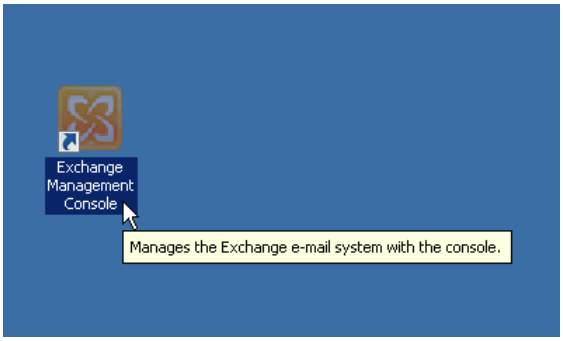

2. Use the **Manage databases** hyperlink or go to **Organization Configuration > Mailbox > Database Management**.

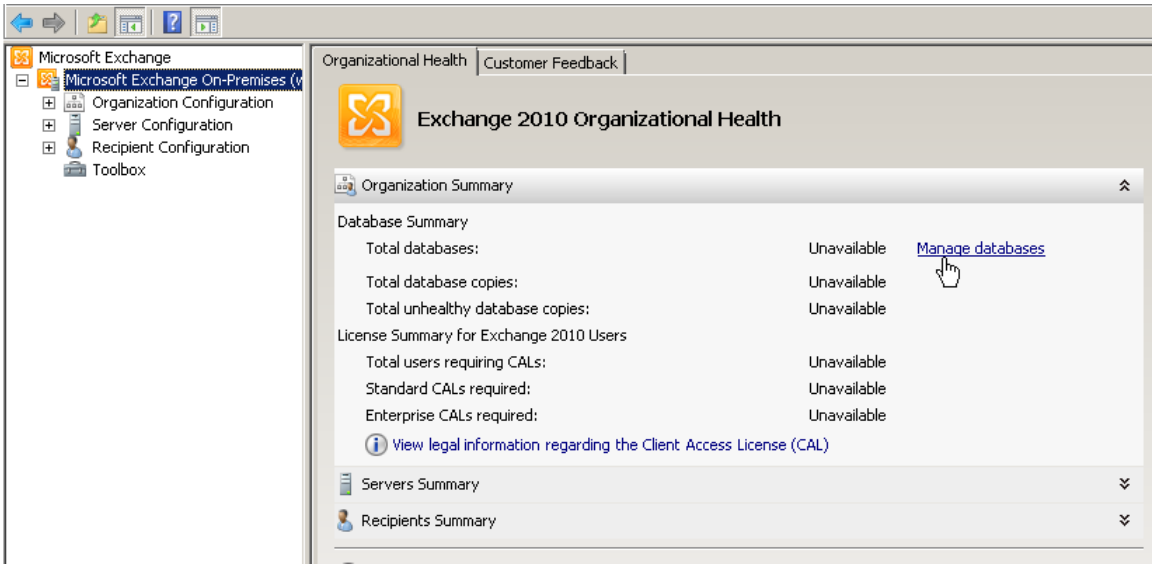

3. Find the created RSG/RDB. Right click on it, then select **Dismount Database**.

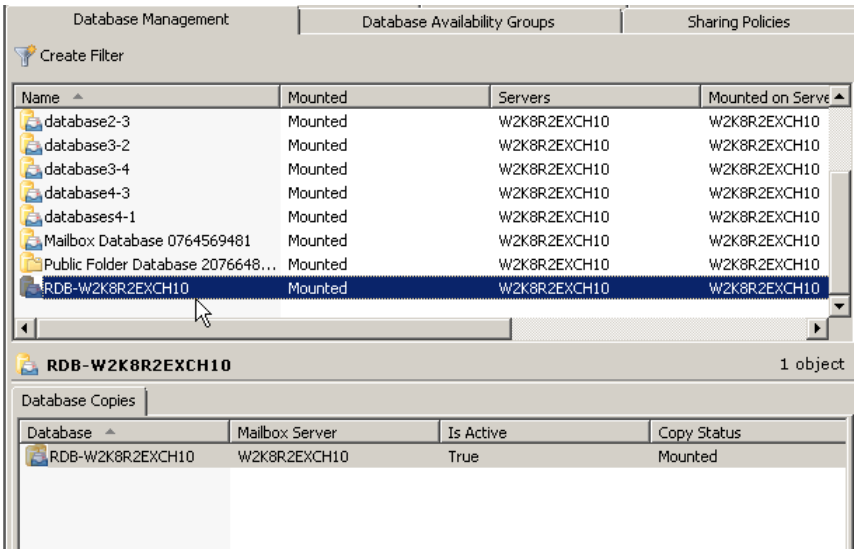

4. Confirm the operation.

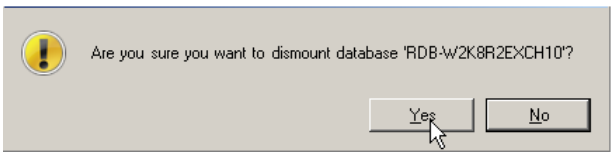

5. When done, either copy the database folder to a machine where MS Outlook with embedded PEGR is installed or share it to access over net.

Installing PEGR

We won't go into details on how to install PEGR here. All details you can get in its help or online manual. Let's just mention a couple of crucial things about it:

- MS Exchange ESE libraries are available* (ESE.dll, ExchMem.dll, Exosal.dll, and JCB.dll for Exchange 2003; ESE.dll, ExchMem.dll, and JCB.dll for Exchange 2007; ESE.dll, Store.exe for Exchange 2010);
- MS Outlook 2003/2007/2010 is installed;
- At least one outlook profile is configured**;

* PEGR can either be installed on a computer with Exchange Server or without it. In the last case you will have to additionally provide ESE libraries during the installation, which you can find in a binary folder of an installed MS Exchange Server (by default, **C:\Program Files\Microsoft\Exchange Server\Bin**). It's not a limitation of PEGR, it's a limitation imposed by Microsoft.

** You're free to configure an outlook profile after the installation is over.

Working with contents of Exchange RSG/RDB in MS Outlook

Once PEGR has been installed, please do the following to restore email items from the created Exchange RSG/RDB:

1. Launch MS Outlook with administrative rights.
2. Click on the **Add an Exchange Database Archive** icon found in the MS Outlook toolbar.

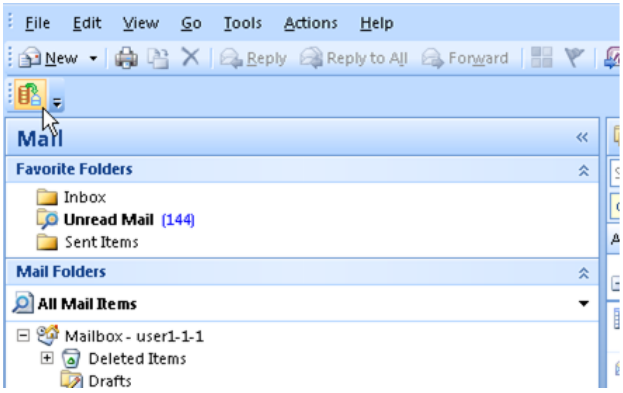

3. Provide your registration info.

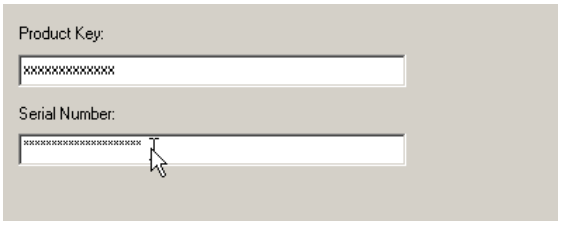

4. Select **Microsoft Exchange Database File.**

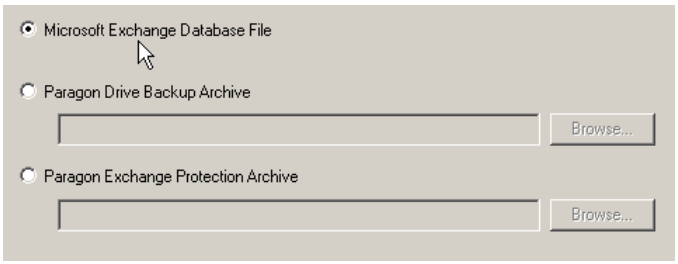

5. Browse for the created Exchange RSG/RDB or set a full path to it manually. Click **Next** to proceed.

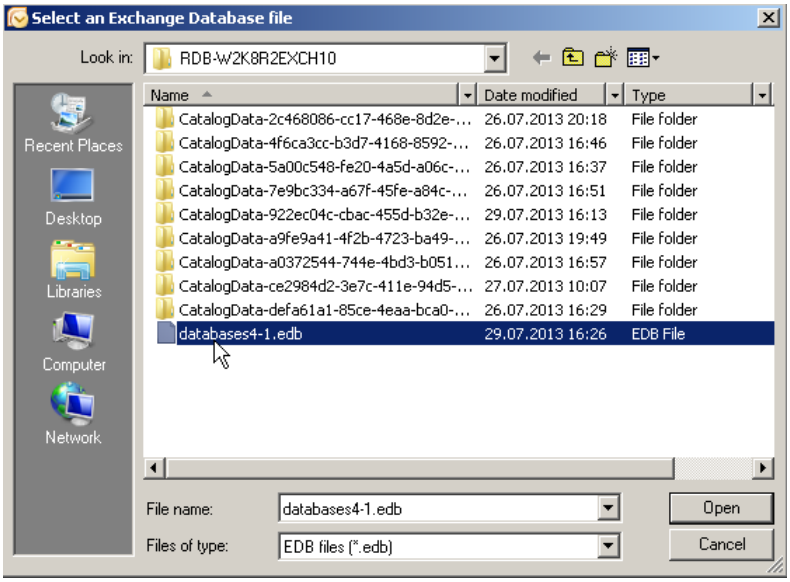

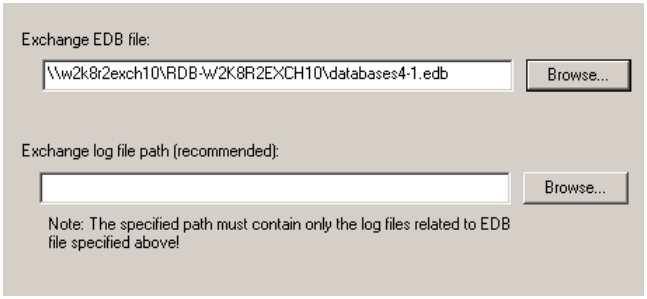

6. Select a mailbox or several mailboxes that contain required email items. Click **Finish** when ready.

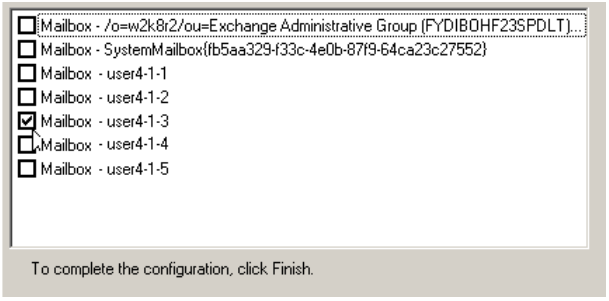

7. As a result you'll get a new mail folder called **Paragon Mail Archive**.

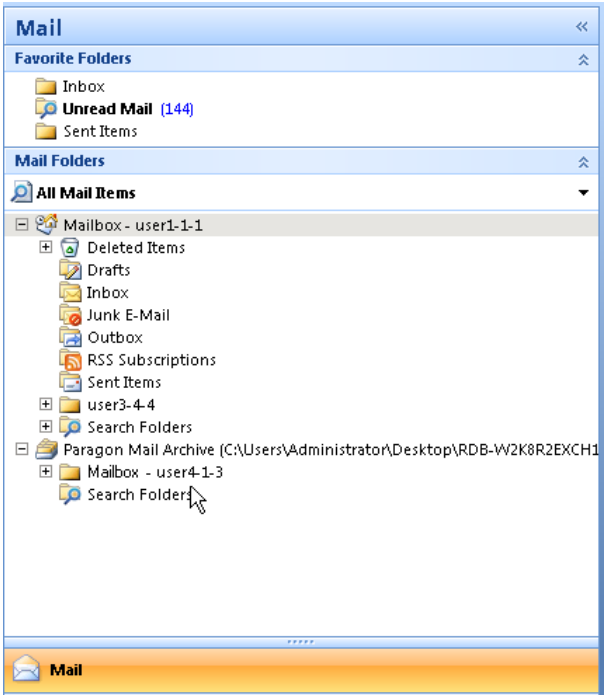

8. Thus you can view contents of the selected mailbox(es), search for and export a single folder, message, contacts item, note, etc. through the standard facilities of MS Outlook.

To know more on the subject, please consult documentation that comes with MS Outlook.

Administering Infrastructure

General Settings

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Logo** button, then select **Settings**.

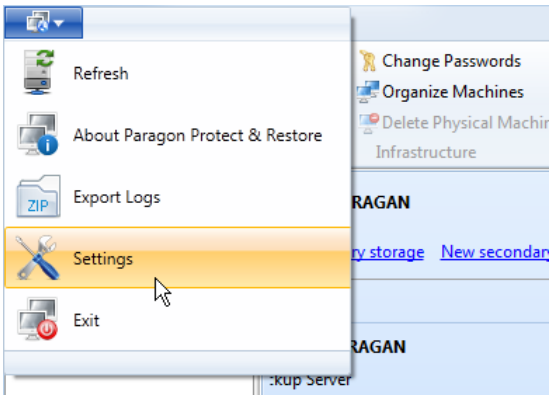

3. All program settings are grouped into several sections: **Console**, **Backup**, **Infrastructure**, and **Accounts & Groups**. By selecting a section from the list in the left pane, you open a set of corresponding options in the right.

Configuring Console

This section contains two subsections that enable to adjust the working environment and the logging level.

Application tab:

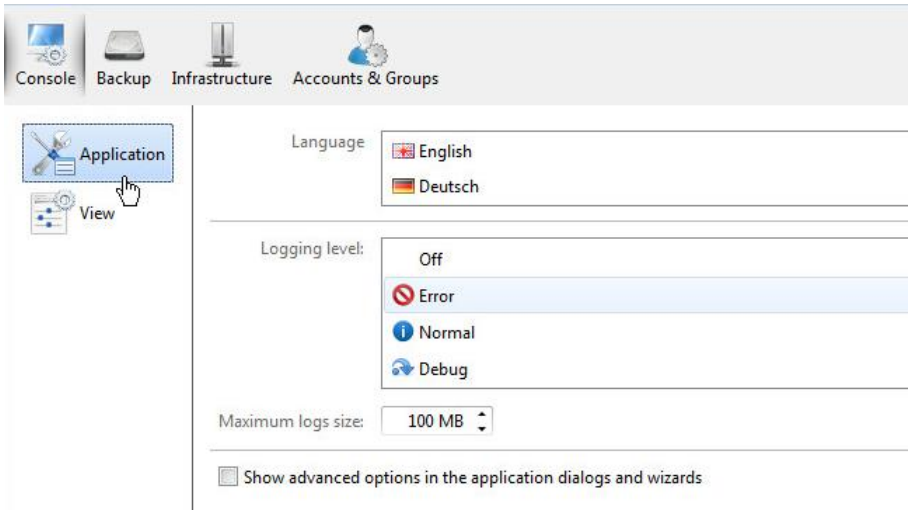

- **Language**. Specify the required language from the list of available interface localizations.
- **Logging level**. By default, there will only be stored console logs generated by unsuccessful (failed) activities (the **Error** option). If you choose **Normal**, logs of any activity will be stored. The **Debug** level is for developers or when cooperating with Paragon's Support Team.
- **Maximum logs size**. By default, only 100 MB of the console log files will be retained.
- **Show advanced options…** Mark this option to enable additional options during creation of storages, VM backup or replication policies, etc.

View tab:

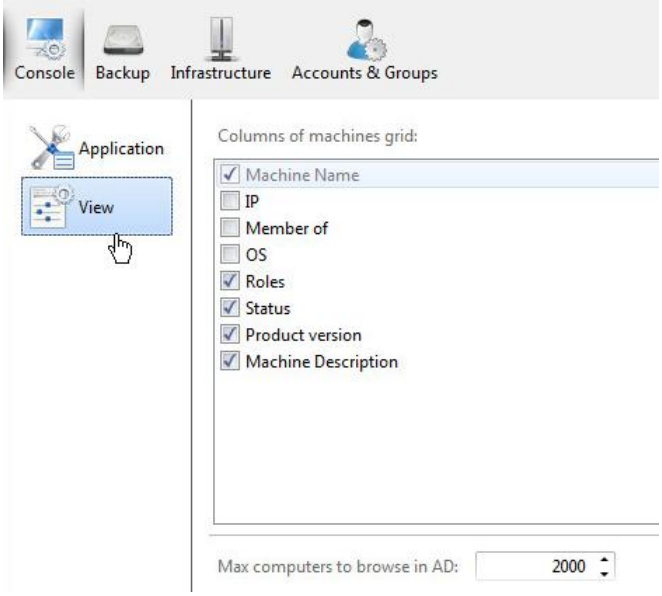

- **Columns of machines grid**. Specify properties for computers that join the infrastructure you'd like to be shown on the Machines panel.
- **Max computers to browse in AD**. If you've got more than 2000 machines in an Organizational Unit, then their browsing and listing in our console may lead to performance degradation. To avoid this, the following limitation has been introduced. Please use the **Find machines** option to find the required machine.

Configuring Backup

In this section you can specify what data should be automatically ignored during physical backup operations. You can filter certain files by creating exclude masks to effectively manage contents of backup images. By default, there are no available filters. To create a filter, please click the **Add new category** hyperlink.

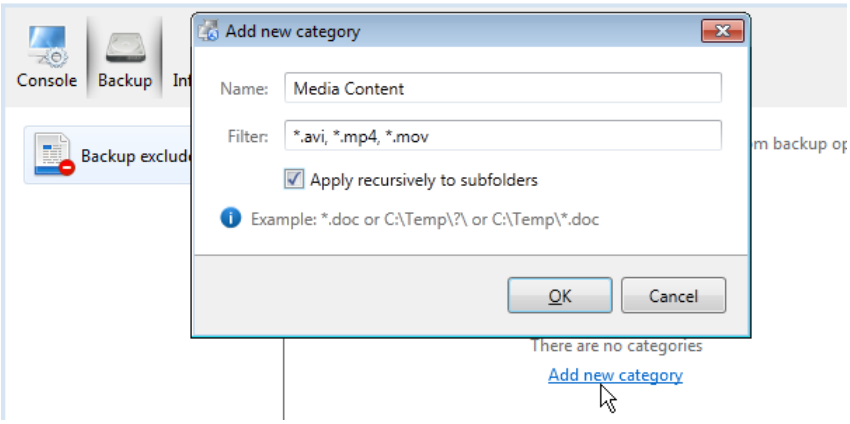

In the opened dialog you should give a name to the filter and specify a filter mask by using * or ? wildcards. Mark the **Apply recursively to subfolders** option to make the program apply the specified masks to all subfolders. By clicking **OK** a new item will appear on the list of filters.

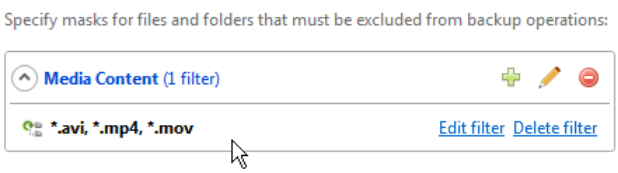

Exclude filters specified in this section will be applied to all physical backup policies. If you'd like to apply a special filter to one backup policy only, please create a private filter during its configuration.

Configuring Infrastructure

This section contains two subsections that enable to configure events and tasks retention policies and the system behavior.

Information retention tab:

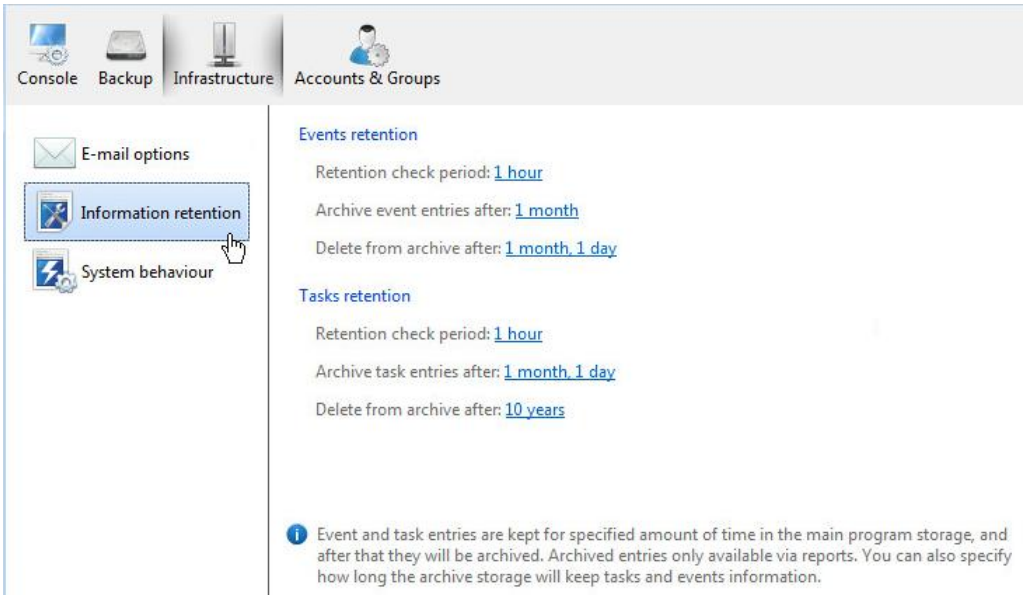

PPR includes a powerful logging mechanism, which enables to keep track of any action accomplished either by the administrator or the infrastructure itself. To avoid overgrowth of the log database, it can be truncated by setting the maximum object lifetime until archived or deleted.

- **Events retention**.
  - *Retention check period*. By default, retention of events stored in the main database of PPR will be accomplished every hour.
  - *Archive event entries after*. By default, events older than 1 week will be automatically archived.
  - *Delete from archive after*. By default, events older than 2 weeks will be automatically deleted.
- **Tasks retention**.
  - *Retention check period*. By default, retention of tasks stored in the main database of PPR will be accomplished every hour.
  - *Archive task entries after*. By default, tasks older than 1 week will be automatically archived.
  - *Delete from archive after*. By default, tasks older than 10 years will be automatically deleted.

System behavior tab:

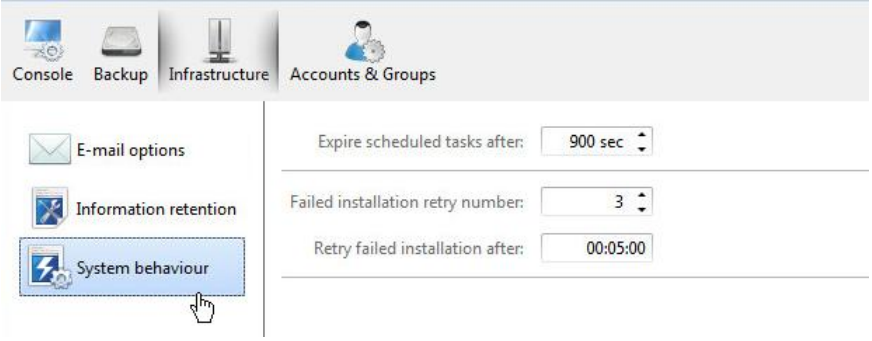

- **Expire scheduled tasks after**. By default, a pending backup task, which hasn't been started at the specified time (local time of each target machine is taken into account) is stated as failed after 900 sec. If setting this value to zero and one or several target machines are offline or temporarily unavailable at the time of backup, the backup policy will fail immediately with a corresponding error. So by this option you can give some credit to your scheduled backup tasks.
- **Failed installation retries**. By default, a failed installation policy will be re-submitted three times pausing between the attempts as specified in the next option. If no success, it will be aborted with a corresponding error. Change the default value if necessary.
- **Retry failed installation after**. By default, a failed installation policy will be re-submitted in five seconds if two or more attempts are specified in the previous option. Change the default value if necessary.

Configuring Security Groups

In this section you can specify user accounts and their privileges in administering the PPR Infrastructure. Several predefined PPR security groups are automatically created during the product installation (locally for a workgroup environment and in Active Directory for a domain):

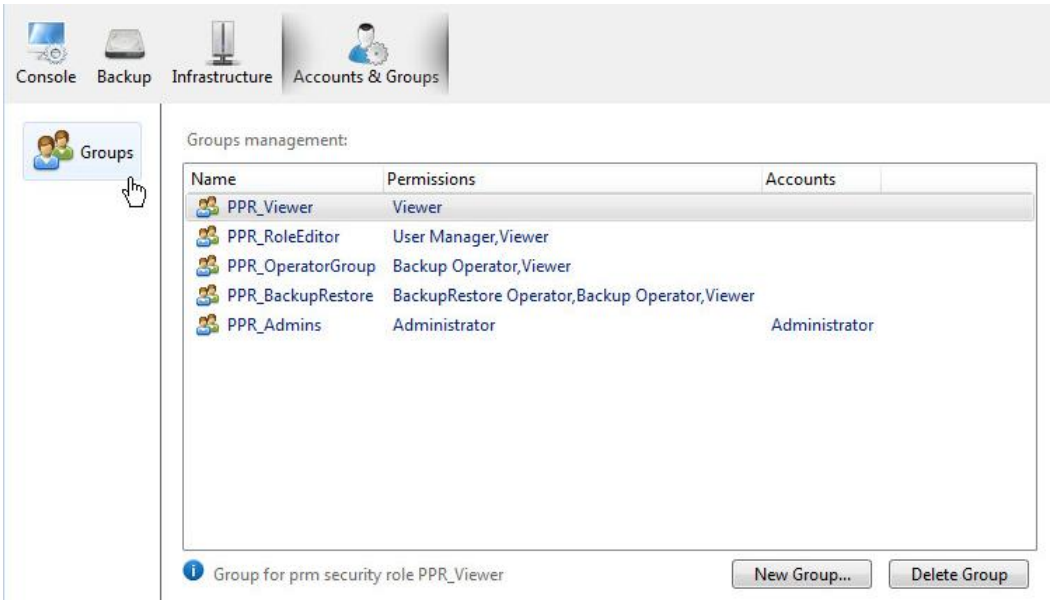

- **PPR_Admins**. A local or domain administrator account you used to install the PPR Server (Administration Server + Installation Server) automatically gets here. Members of this group have all possible permissions. Click here for more information.
- **PPR_BackupRestore**. Members of this group are allowed to initiate backup and restore operations and carry out some other operations. Click here for more information.
- **PPR_RoleEditor**. Members of this group are allowed to add/remove members in any group and carry out some other operations. Click here for more information.

- **PPR_Operator**. Members of this group are allowed to initiate backup operations and carry out some other operations. Click here for more information.
- **PPR_Viewer**. Members of this group are only allowed to view activities of the PPR Infrastructure. Click here for more information.

As you see all PPR security groups have a unique prefix "PPR_", which helps to avoid collisions with existing local or AD groups.

The number of operations available in the PPR Console depends on the user account you're logged in.

By using the corresponding buttons you can create a new group or remove an existing one.

Editing group properties

1. Double click a group to see all group members and their permissions in the opened dialog. Beside other parameters, you're allowed to change the number of available permissions as well as add a new member or remove an existing one. Click **Add…** to add a new member to the group.

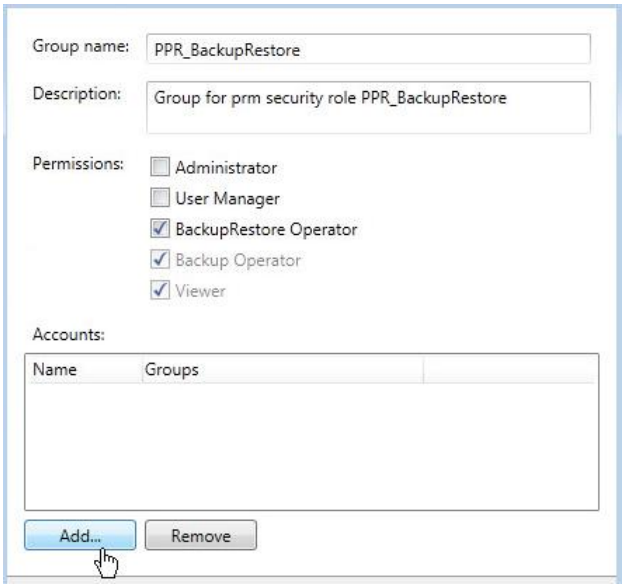

3. Either choose an existing user account from the list or create a new one by providing a user name and password. Please note that all password properties (number of characters, complexity, expire date, etc.) are imposed and administered by means of Windows OS. Click **Add** when ready.

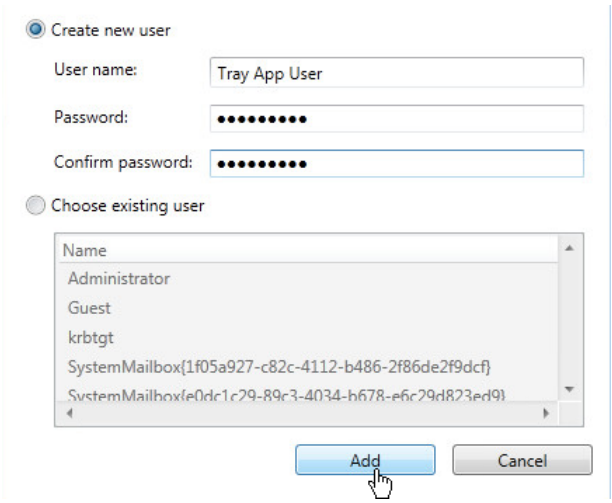

4. The specified user account should get to the list of group members.

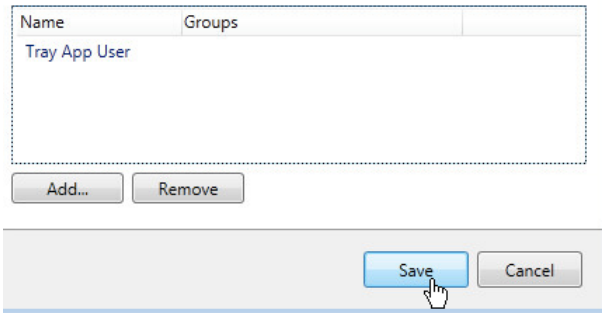

We strongly recommend you to modify properties of existing PPR security groups through the PPR Console. Anyway if you did it directly in AD for instance, please additionally introduce the same changes in Prm.Common.Service.exe.Config, which you can find here: C:\Program Files\Paragon Software\Remote Management\program

Permissions of security groups

Permissions	Administrator	User Manager	BackupRestore Operator	Backup Operator	Viewer
Edit application settings	+	+	Add a recipient address only	-	-
Build WinPE Recovery Media	+	-	+	-	-
Manage Machines					
Add/Remove/Update/Upgrade	+	-	-	-	-
Export logs	+	-	+	-	-
Manage Backup Server					
Add/Remove	+	-	-	-	-

Browse	+	+	+	+	+
Manage Backup Storage					
Set up archiving/retention/deduplication	+	-	-	-	-
Rebuild storage	+	-	-	-	-
Set up free space notifications	+	-	+	-	-
Advanced settings	+	-	-	-	-
Manage Backup Data					
Replica failover/test failover	+	-	+	-	-
Launch backup	+	-	+	-	-
Export backup	+	-	+	-	-
Delete backup	+	-	+	-	-
Check backup Integrity	+	-	+	-	-
Repeat archiving	+	-	-	-	-
View archiving history	+	+	+	+	+
Backup/Restore Policies					
Create/Modify/Delete	+	-	+	-	-
Initiate/Cancel/Restart	+	-	+	Backup only	-
Maintenance Policies					
Run/Restart/Cancel/Disable archiving	+	-	+	-	-
Delete/Modify archiving	+	-	-	-	-
Run/Restart/Cancel/Disable retention	+	-	+	-	-
Delete/Modify retention	+	-	-	-	-
Manage ESX Connections					
Add/Modify/Delete	+	-	+	-	-

Configuring Notifications

In this section you can configure so called notification collectors – in fact these are endpoints where you would like generated notifications to be sent by PPR. Our product supports four types of collectors: File, SMTP, SysLog and WinLog (Windows Event Log).

To set up a notification collector, please do the following:

1. Select the required collector type on the left, then click the **Create a new collector** link.

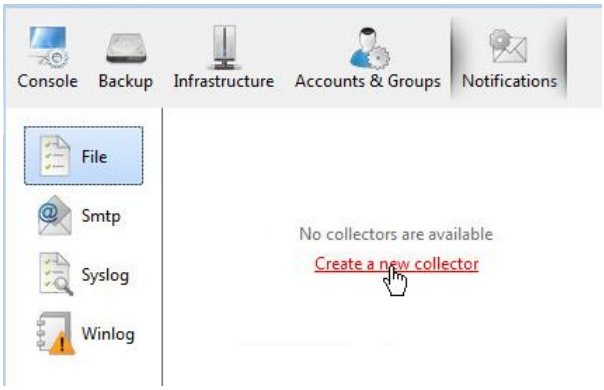

File collector:

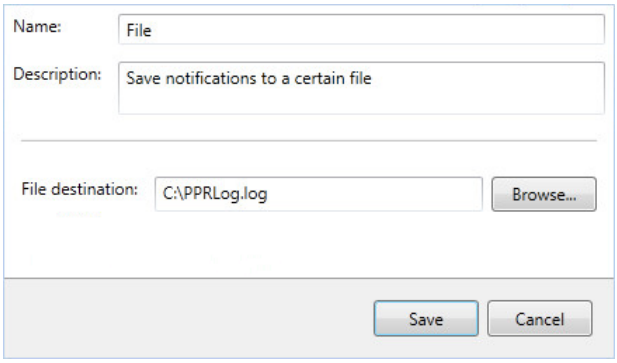

- **Name**. Give a unique name to the created collector to differentiate it from the others.
- **Description**. Provide a detailed description (optional).
- **File destination**. Browse for a folder, then modify the default file name if necessary. Alternatively, you can set a full path manually by typing it in the corresponding field.

SMTP collector:

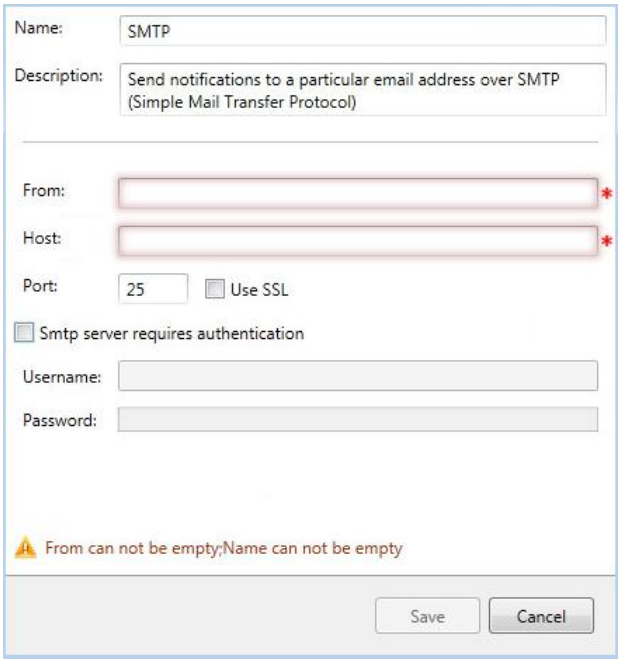

- **Name**. Give a unique name to the created collector to differentiate it from the others.
- **Description**. Provide a detailed description (optional).
- **From**. Provide a name of the sender.
- **Host**. To send email notifications, it is necessary to have access to a computer running an SMTP (Simple Mail Transfer Protocol) server. All outgoing messages are first sent to the SMTP server, which in its turn delivers them to the required recipients. The address may be represented as a traditional Internet host name (e.g.: mail.com) or as an IP numeric address (e.g. xxx.xxx.xxx.xx).
- **Port**. Change the default port if necessary.
- **Use SSL** (Secure Socket Layer). Activate the option to establish a secure connection to the email server.
- **SMTP server requires authentication**. Activate the option to allow the program to authenticate on the SMTP server before sending messages, then provide valid user credentials in the corresponding fields.

Email subject and recipients or addresses where notifications should be sent to are specified during configuration of a notification policy.

SysLog collector:

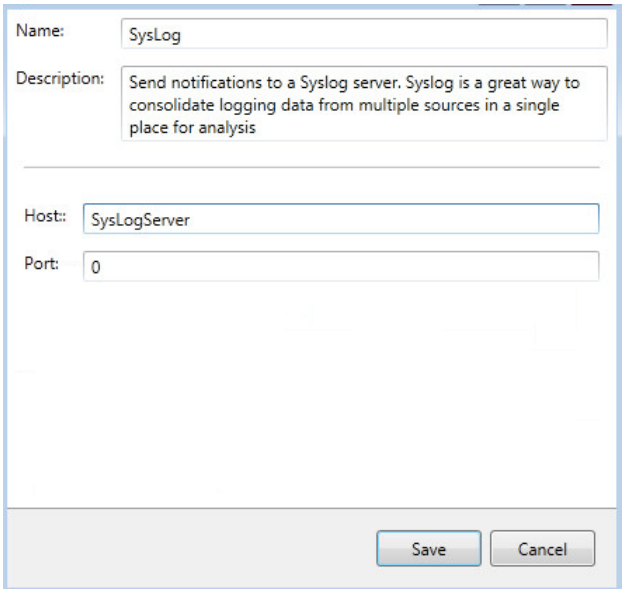

- **Name**. Give a unique name to the created collector to differentiate it from the others.
- **Description**. Provide a detailed description (optional).
- **Host**. Provide a DNS name or IP address of a machine hosting a SysLog server.
- **Port**. Usually, a SysLog server listens to port 514 or 6514 for incoming notifications generated by remote SysLog clients.

WinLog collector:

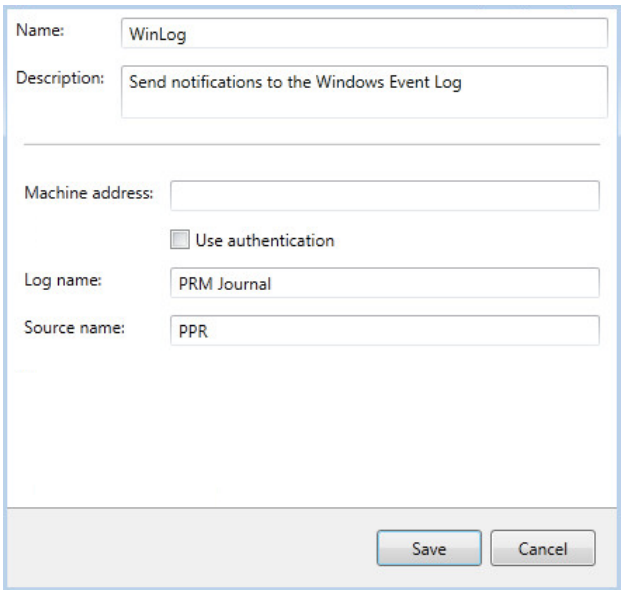

- **Name**. Give a unique name to the created collector to differentiate it from the others.
- **Description**. Provide a detailed description (optional).
- **Machine address**. Provide a DNS name or IP address of a machine acting as Windows Event Collector. Use a corresponding option to provide access credentials if necessary.
- **Log name**. By default, all notifications generated by PPR will be stored in a log file named 'PRM Journal', which you can change as required.
- **Source name**. If having several PPR infrastructures you can use this option to easily pinpoint notifications coming from a particular infrastructure.

2. Click **Save** when ready with all parameters. As a result a new item will appear on the list. By using the corresponding buttons you can create a new collector or modify/delete an existing one.

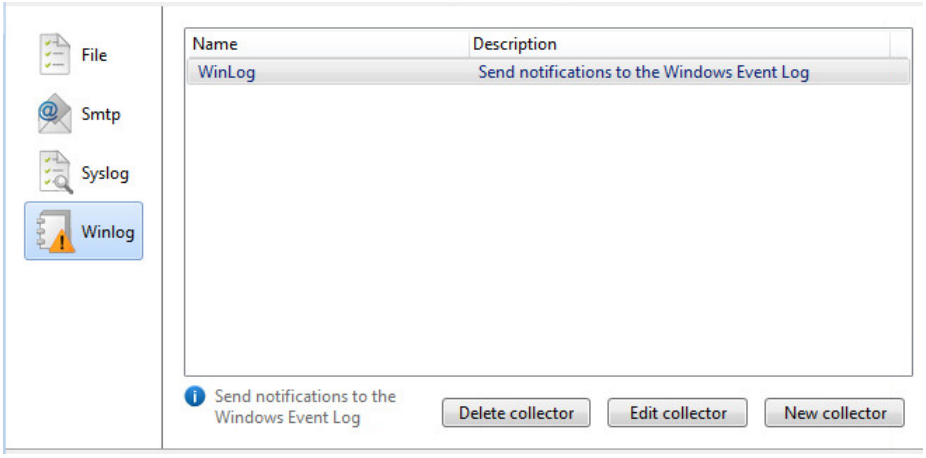

Changing Roles

You can extend or limit the functionality of any member of the infrastructure by adding/removing roles.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon then select **Edit Roles**.

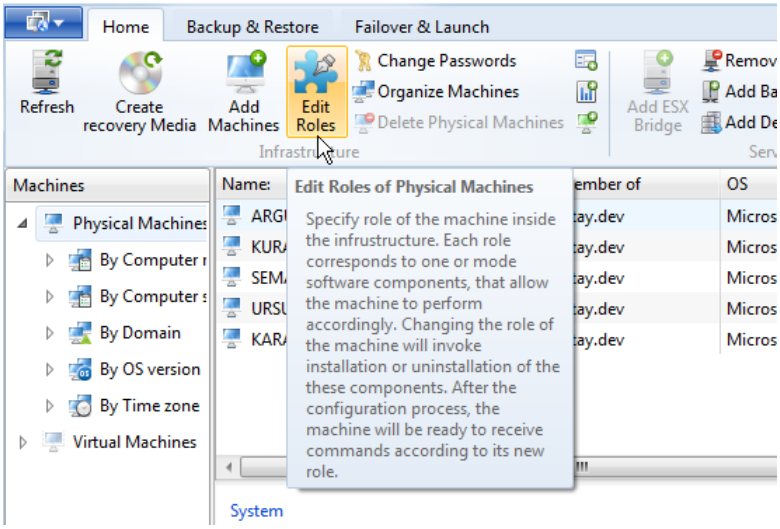

You can also initiate this operation by the right click of the mouse button on the required infrastructure member, then selecting the corresponding option. In this case you won't need to specify a machine you'd like to work with.

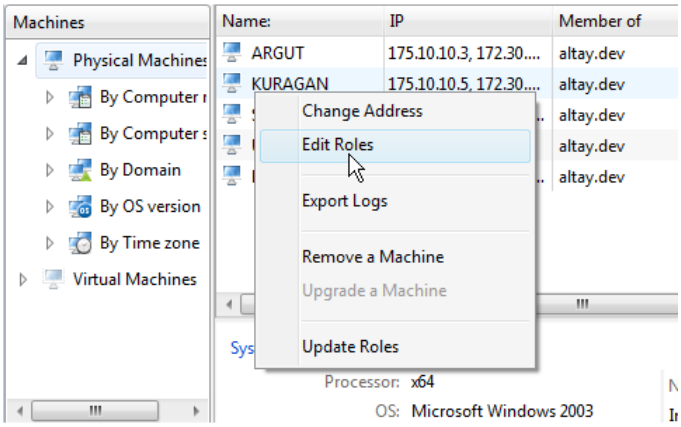

3. The wizard will list all machines that join the infrastructure (have acquired at least the general role of Agent). Mark checkboxes next to those that require plug-in modification. Click **Next** to proceed.

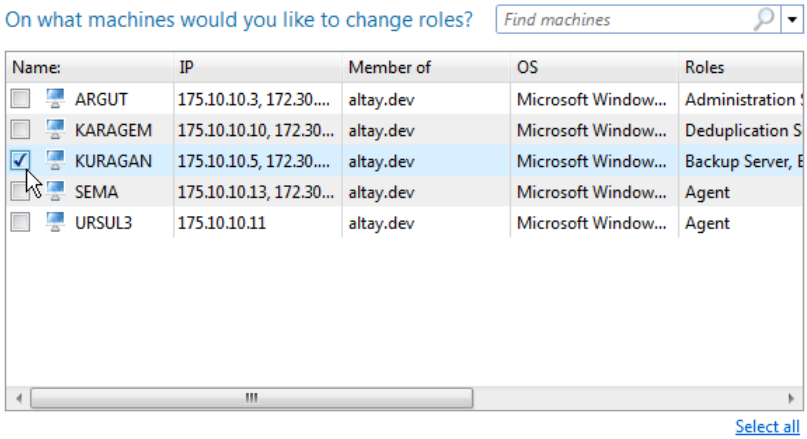

If you'd like a particular set of roles for each machine, please specify one machine at this stage. Otherwise, all selected machines will acquire the same set of plug-ins. When completed, repeat the action for another machine.

4. Specify roles to install.

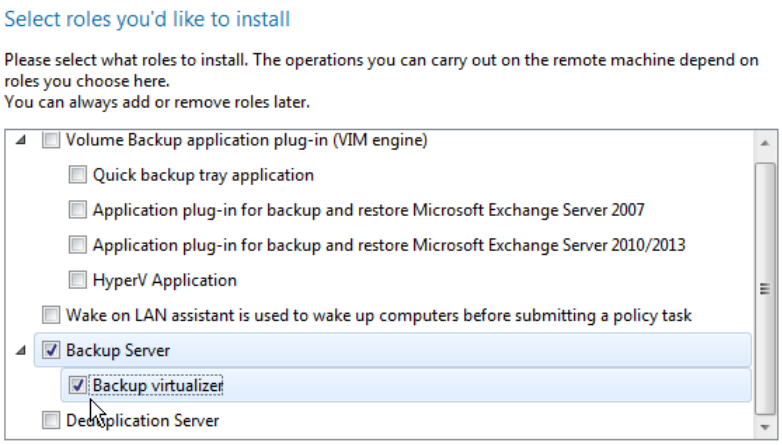

The number of available plug-ins depends on the purchased license.

5. Specify how and when the selected roles should be installed.

Changing Machine Address

If you've renamed a machine that joins the PPR infrastructure, please use this dialog to assign the new name to the corresponding infrastructure member, otherwise our product won't be able to work with this machine.

This operation will be available if the advanced mode is enabled in the Settings dialog.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, select **Machines > Physical Machines**, right click the required infrastructure member, then select **Change Address**.

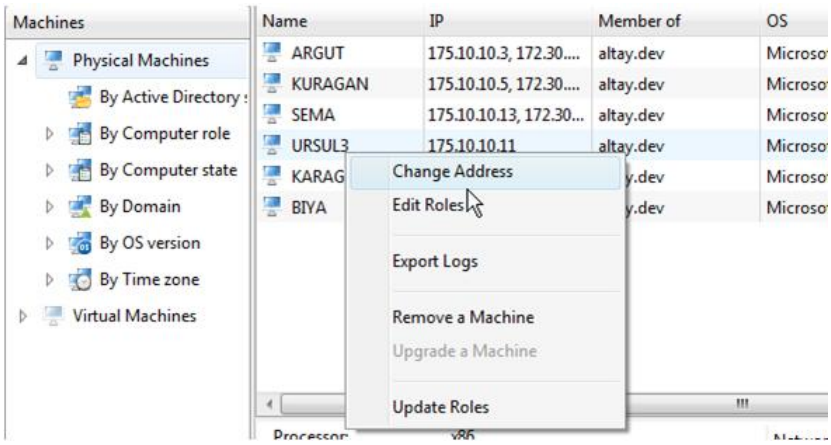

You can also initiate this operation by clicking the **Edit** link opposite name of the currently selected infrastructure member.

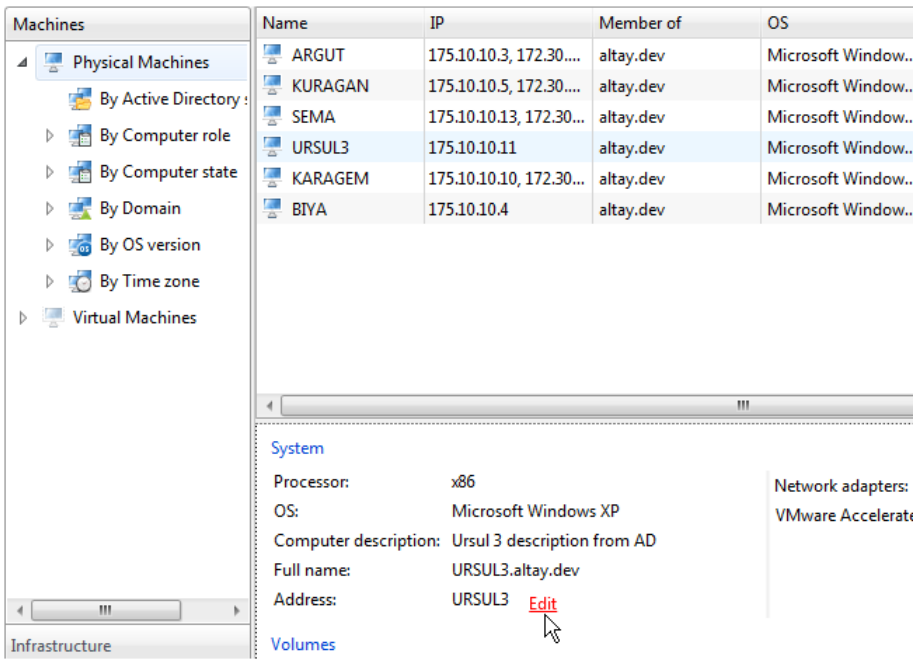

3. Type in the desired name, then click **OK** to confirm.

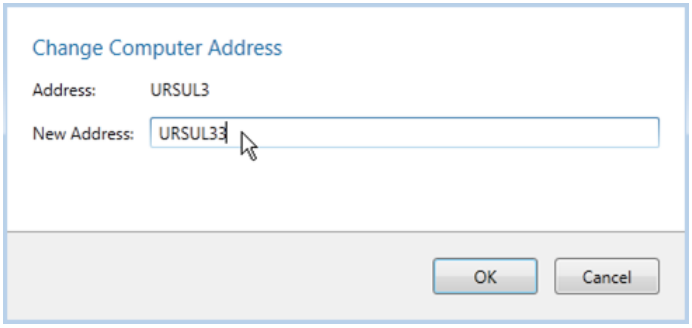

Managing Access Credentials

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon then select **Change Passwords**.

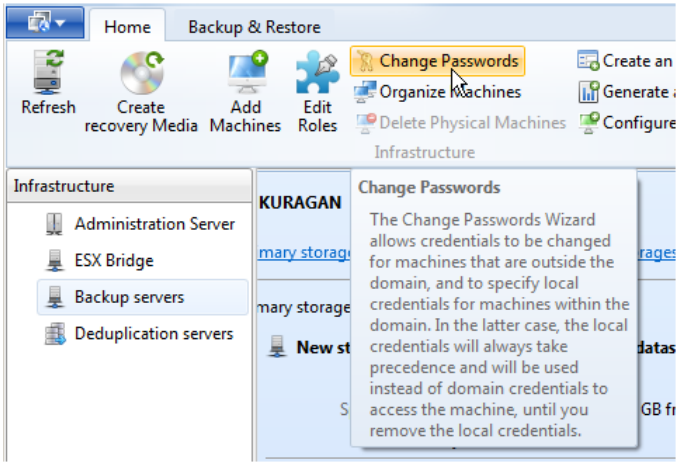

3. The Change Password Wizard will prompt you to choose one of the three options:

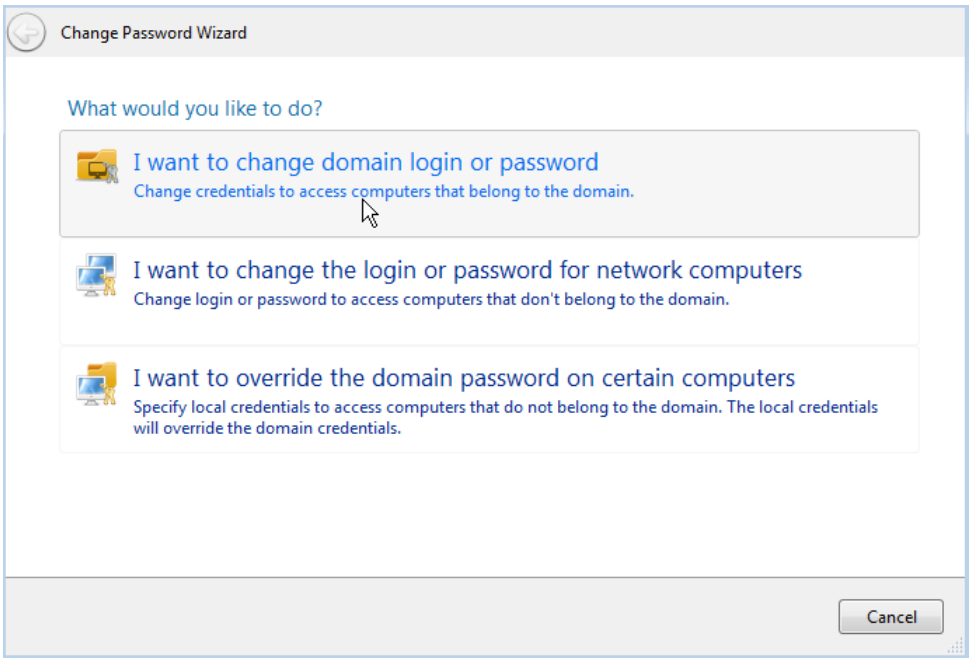

- **Change domain credentials**. If credentials of the domain administrator have been changed due to a planned security routine or accidentally, please use this option to provide these new credentials to our product, otherwise the entire infrastructure won't be able to operate.

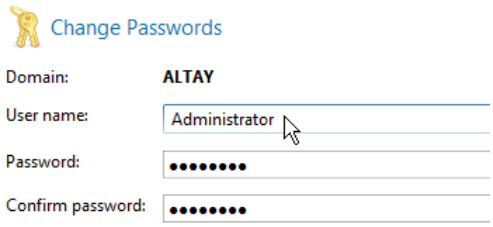

- **Change credentials of non-domain machines**. If credentials used to access a workgroup machine have been changed, please use this option to provide these new credentials to our product, otherwise it won't be able to work with this machine.

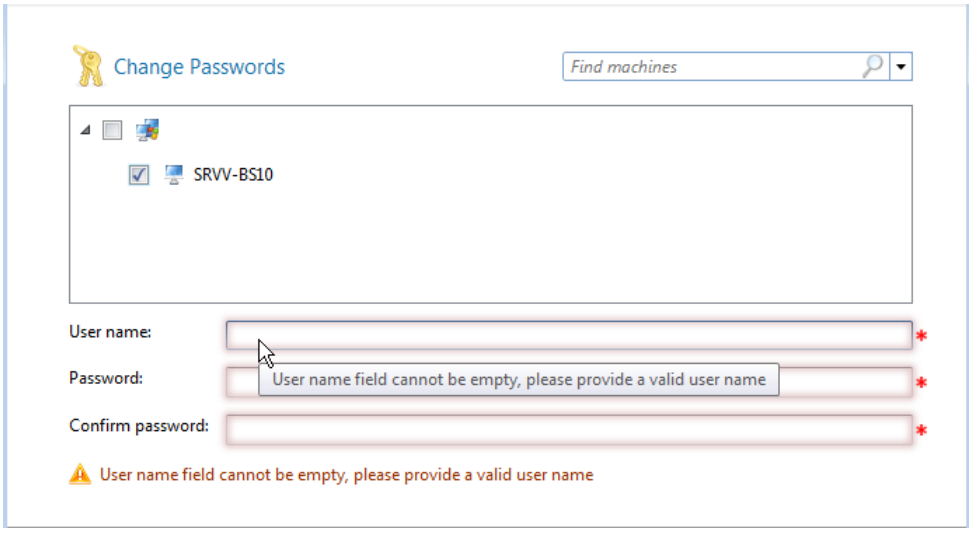

- **Specify local credentials to access particular domain machines**. If you use local system credentials to access one of the domain members of the PPR infrastructure, please use this option to change or remove these credentials, or specify local credentials to any of the domain machine that joins the infrastructure.

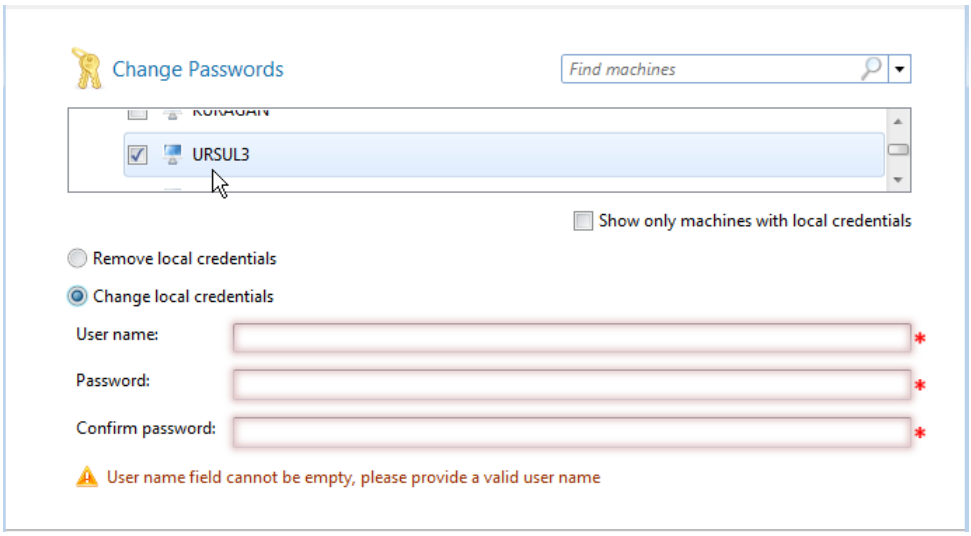

Local credentials set at this step will be used instead of domain credentials, until you remove them.

Managing Policies

Through the **Policies** pane you can administer any of created earlier policies.

Operation scenario

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, go to **Policies**, then click on the required type to see a list of corresponding policies. Use the arrow button below a policy name to see its main properties.

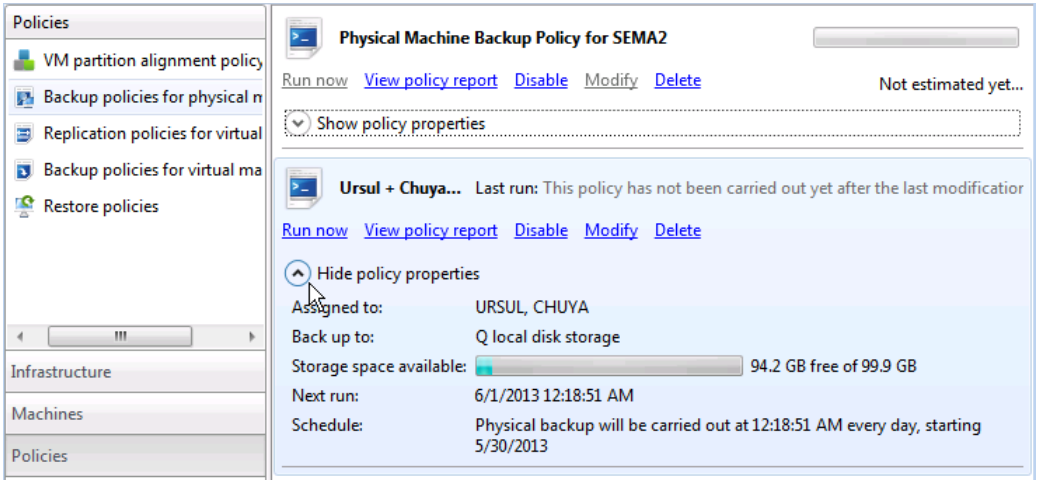

3. **Run now**. Use this option to force launch of the required policy.
4. **View policy report**. Use this option to get detailed information on all policy launches of the required policy. Here you can see when and with what result (succeeded or failed) each policy launch completed. Click **More...** next to a failed policy launch to see the reason.

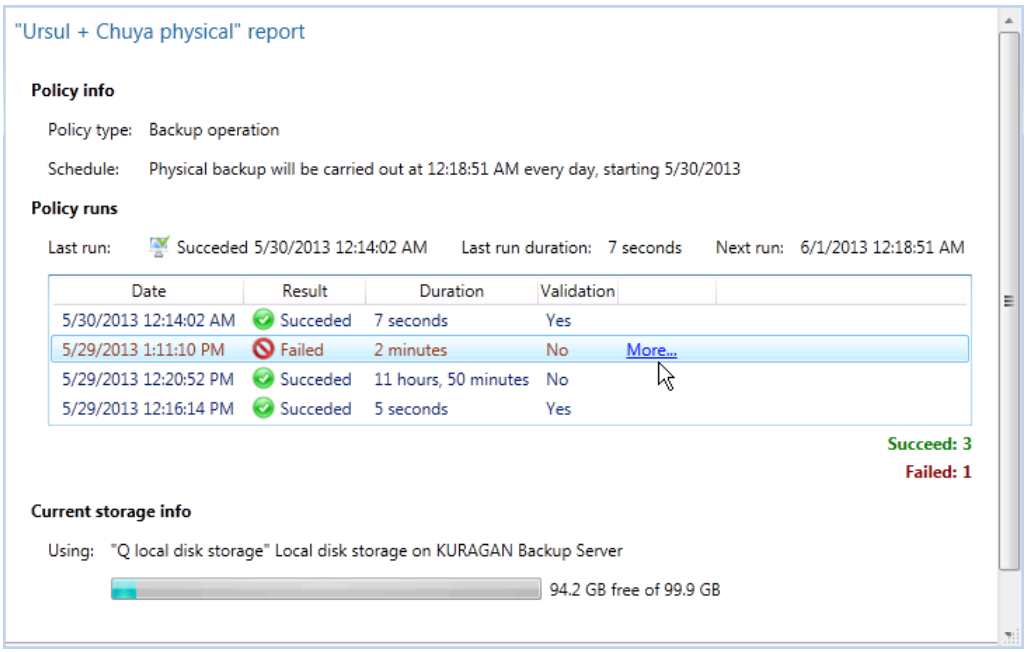

Select a failed object to see error details. If you need more information on the subject, please use the **Technical information** button to see infrastructure logs generated during the specified policy launch.

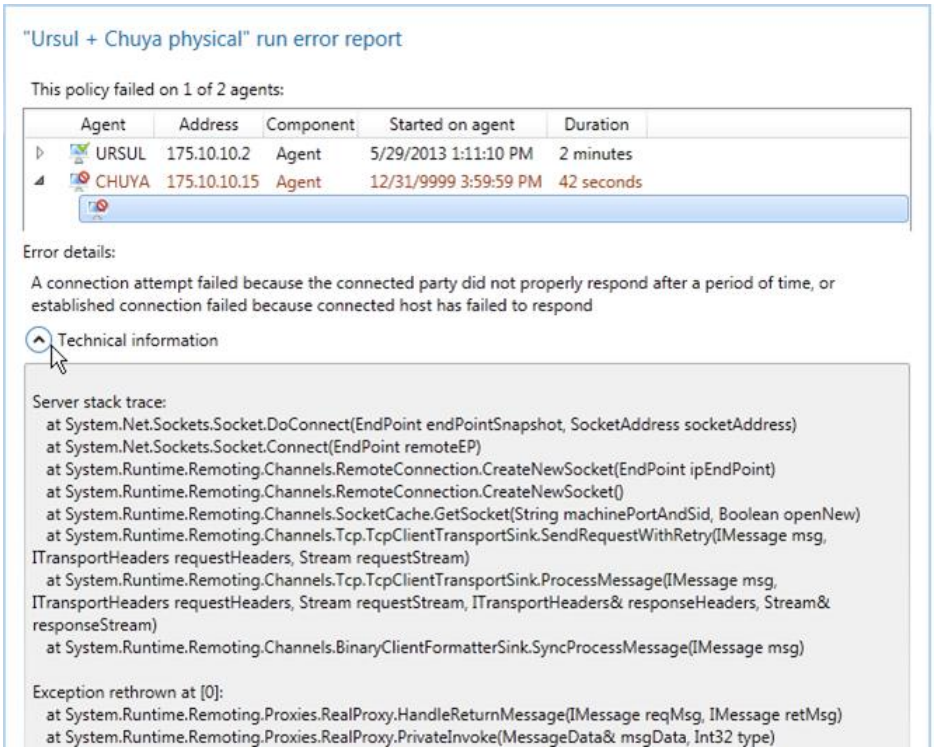

5. **Disable**. Use this option to stop the required policy.
6. **Modify**. Use this option to change properties of the required policy. Please note the number of available properties will depend on the policy type (virtual backup, virtual replication, physical backup, etc.). To know more on the subject, please consult the corresponding chapters: Protecting Virtual Machines, Protecting Physical Machines.

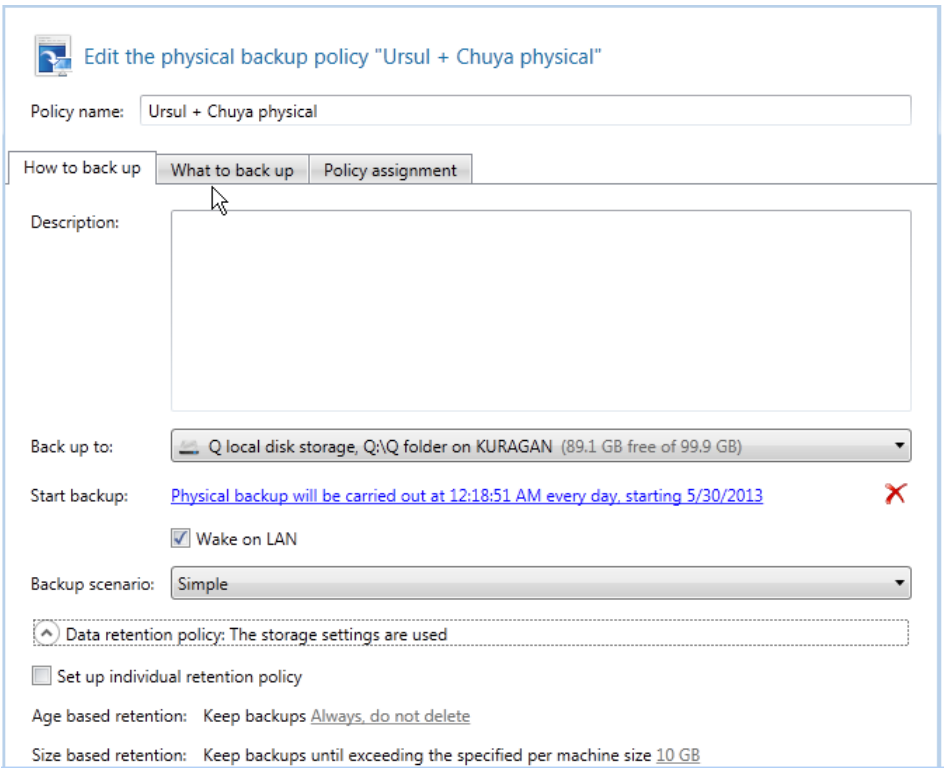

7. **Delete**. Use this option to delete the required policy.

Managing Activities

The **Activities** pane is the primary tool for getting information on policies created by the user. To open the pane, please go to **Activities and Events > Activities**.

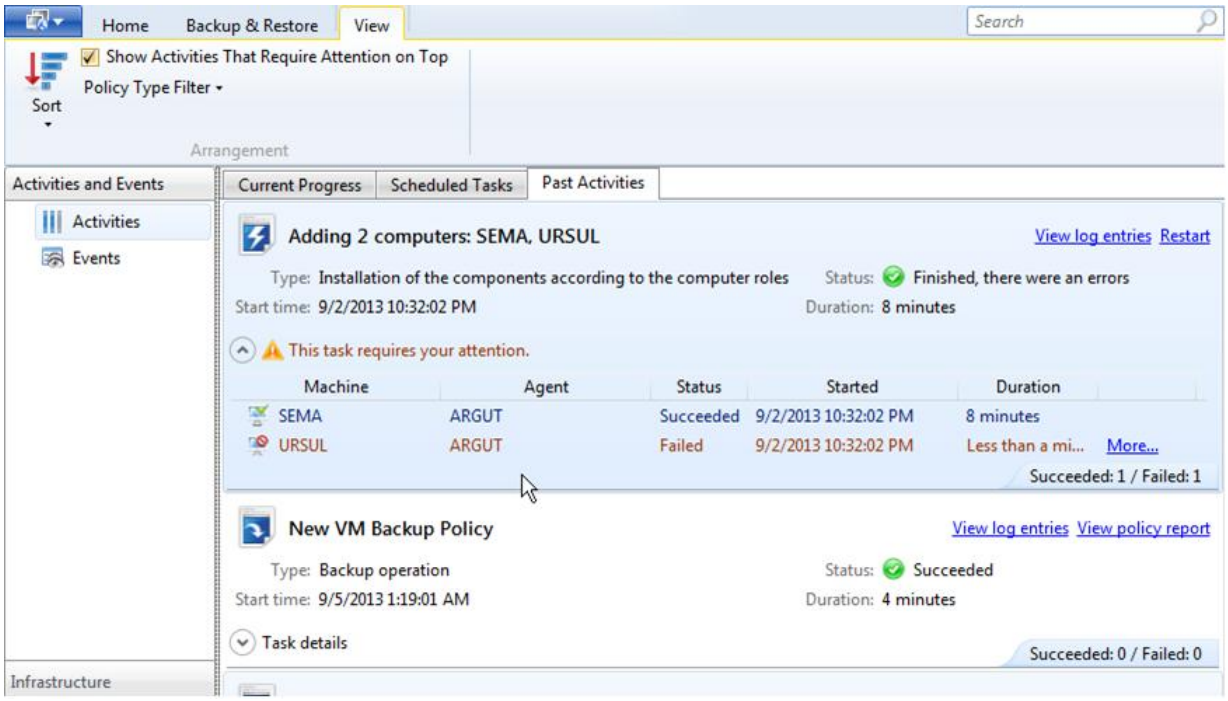

As you can see the pane is divided into three sections:

- **Current Progress** that displays policies being accomplished at the moment,
- **Scheduled Tasks** that displays scheduled policies,
- **Past Activities** that displays already completed policies.

By default, the program lists all ever submitted user policies by name, which can be customized through the **View** ribbon. To make the job with activities easier and more efficient, you can:

- Sort out policies by name, last run, completion time, or duration;

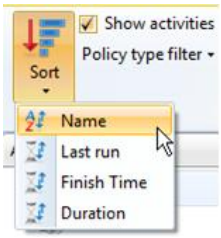

- Filter policies by their type (Installation, VMBackup, VMReplication, PhysicalRecovery, etc.);

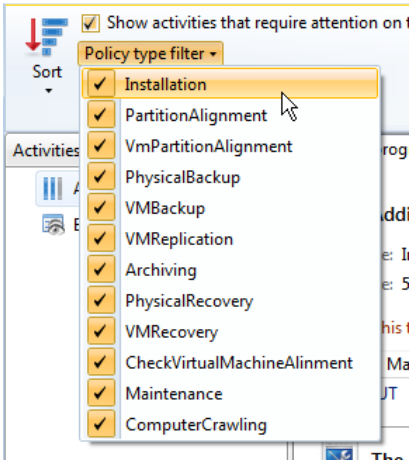

- Make the program display policies that require user's attention on top (partially incomplete or failed policies);

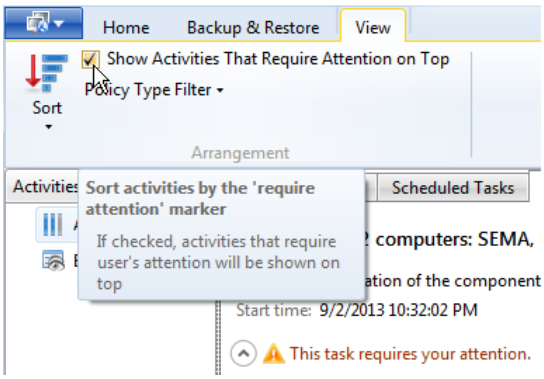

By default, only last run of each policy is displayed. If you'd like to see all launches of the required policy, please use the View policy report option.

Monitoring running activities

Let's take a closer look at one of the running activities to see what information you can get from it and how it can help.

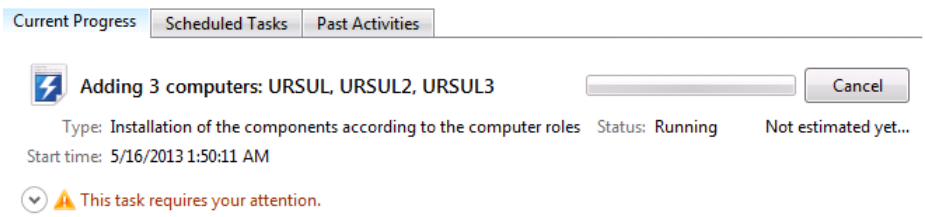

For each activity the program outputs its name and type, time when it's started, its current status, and estimated duration. As you may have noticed the current status of the questioned activity is 'Running' – actually nothing to worry about, if not for a warning that should attract our attention. By clicking on this warning we learn that the installation policy was accomplished by the machine named ARGUT (Protect & Restore Server) with the resulted status 'Finished, there were errors' – not a good sign.

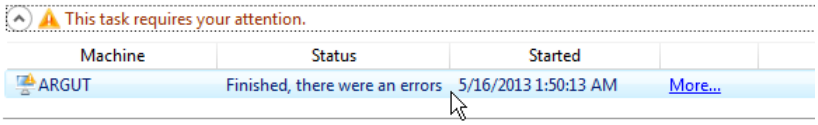

By clicking **More…**, we can finally see the reason.

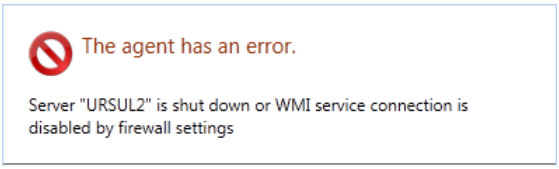

Having this information at hand, you can check the target machine and solve the problem. When done, the installation will be automatically completed.

Monitoring past activities

Submitted policies can either be successfully completed or with errors, which can be seen from their statuses. For each past activity you can open a list of corresponding infrastructure events to get detailed information on the operation progress.

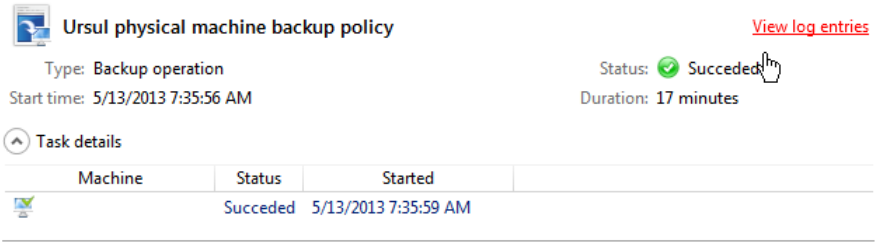

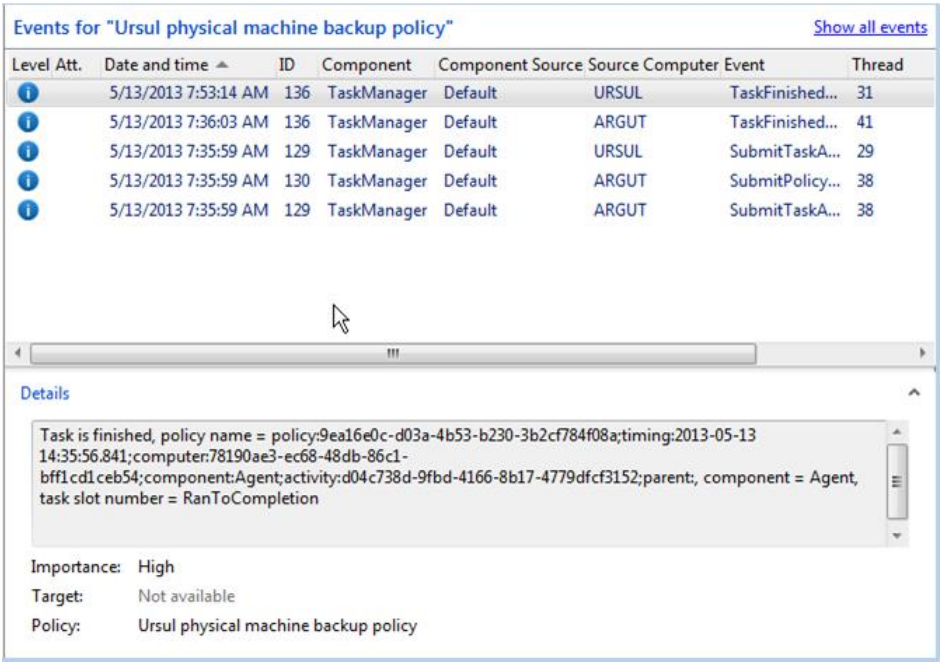

Failed policies can be examined and restarted.

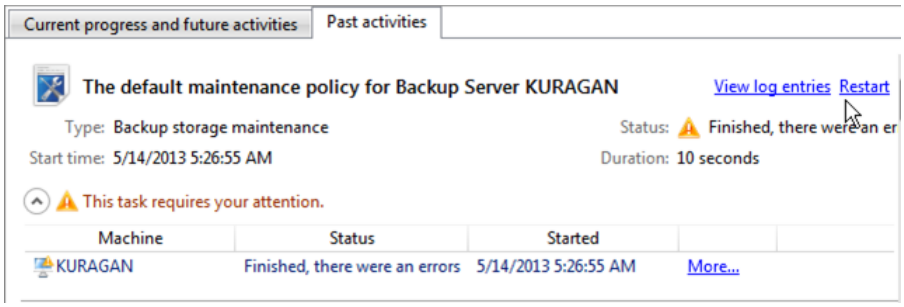

Monitoring scheduled activities

For each scheduled policy you can see when it's planned to run next time. By using corresponding options, it's also possible to change parameters of the policy or to see a detailed report on all its launches.

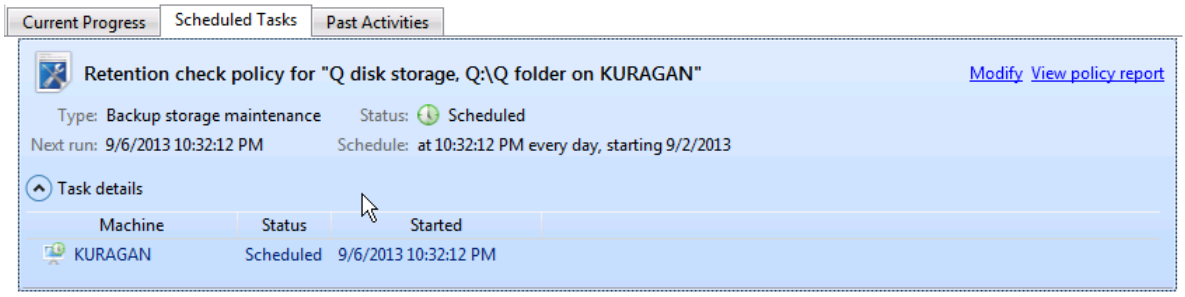

Managing Events

The **Events** pane provides in-depth information on all actions that take place within the infrastructure, both user and service. It's a tool for thorough analysis of the infrastructure functioning and troubleshooting. To open the pane, please go to **Activities and Events > Events**.

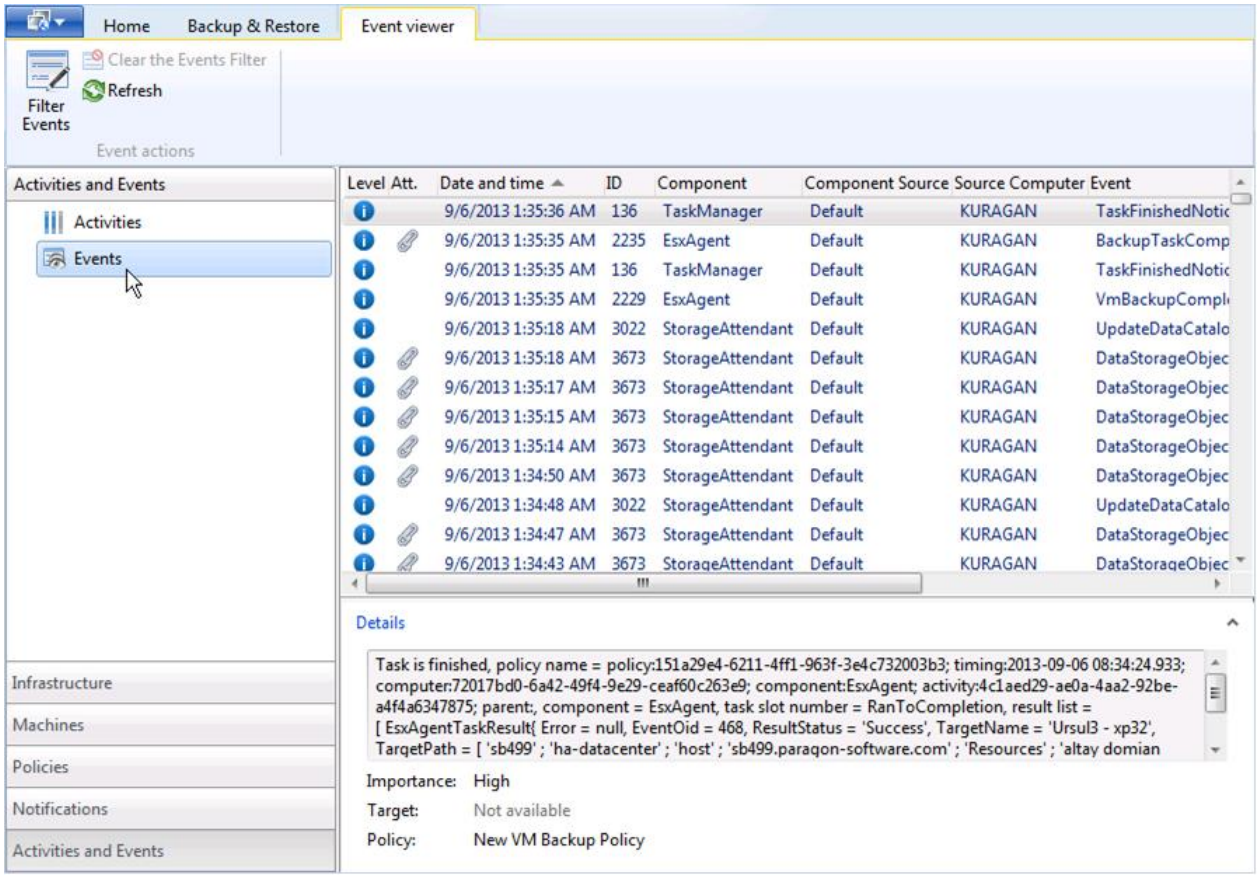

By clicking the arrow button you can hide/unhide details on the selected at the moment event.

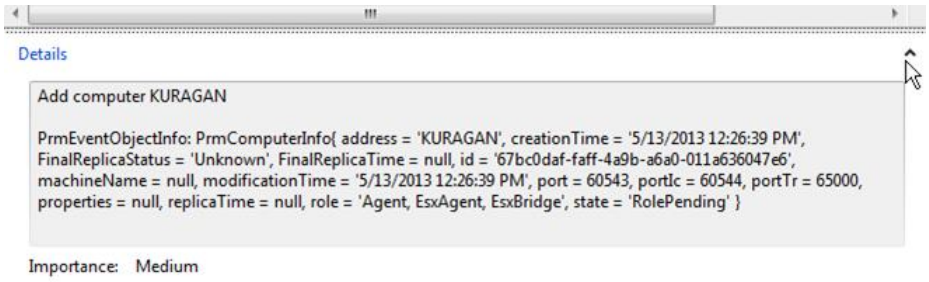

Events can be filtered for easier monitoring if necessary by using the corresponding icon.

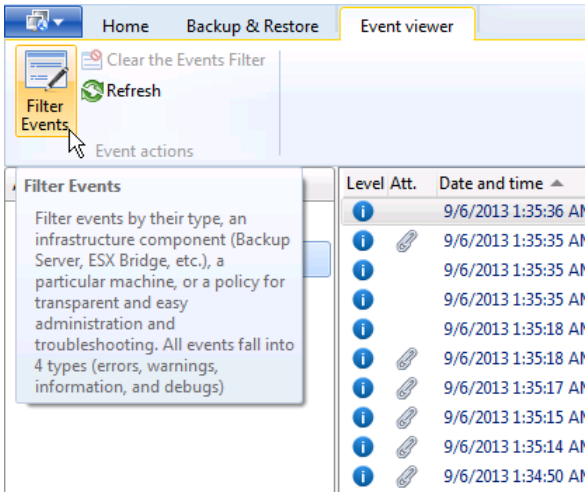

In the opened dialog you can filter events by their type, an infrastructure component (Backup Server, ESX Agent, etc.), a particular machine, or a policy for transparent and easy administration and troubleshooting. All events fall into 4 types (errors, warnings, information, and debugs).

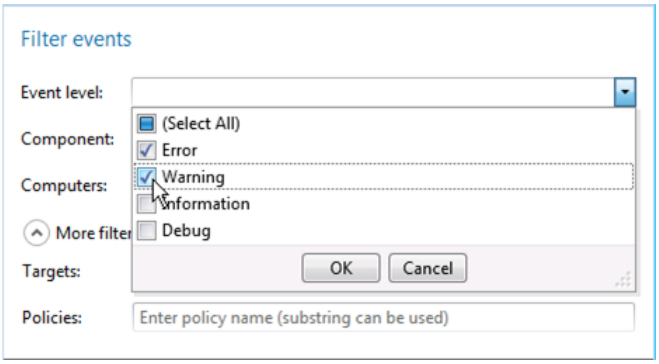

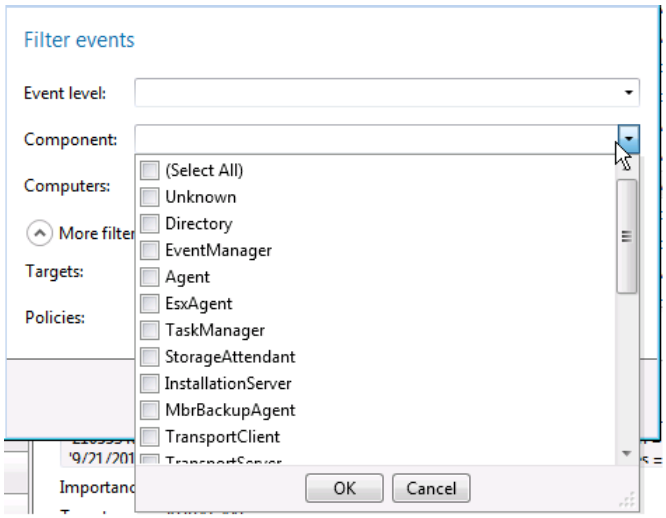

To get the latest events available, force the update by directly requesting Administration Server. This action can take some time.

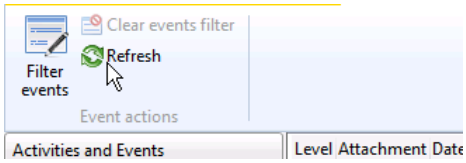

To see all infrastructure events just clear the previously made filtering.

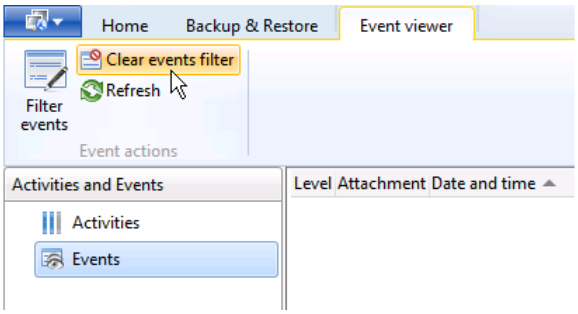

Notifications

PPR includes a powerful tool of monitoring events within the infrastructure from outside called **Notifications**. Infrastructure events you can be notified about are logically divided into two groups. In short, the first group deals with tasks derived from user policies (backup, restore, failover, archiving, etc.), while the second covers all infrastructure events, including system policies, login/logout statistics, etc.

A handy wizard helps you to configure notification policies with minimal effort. It includes ready-made notification templates for most user policies and infrastructure events you may be interested in. For advanced users there's the option to create templates from scratch or modify existing ones. Generated notifications can be saved in a file, the Windows Event Log (WinLog), sent to an email address or a SysLog server (SysLog).

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon, then select **Configure notifications**,

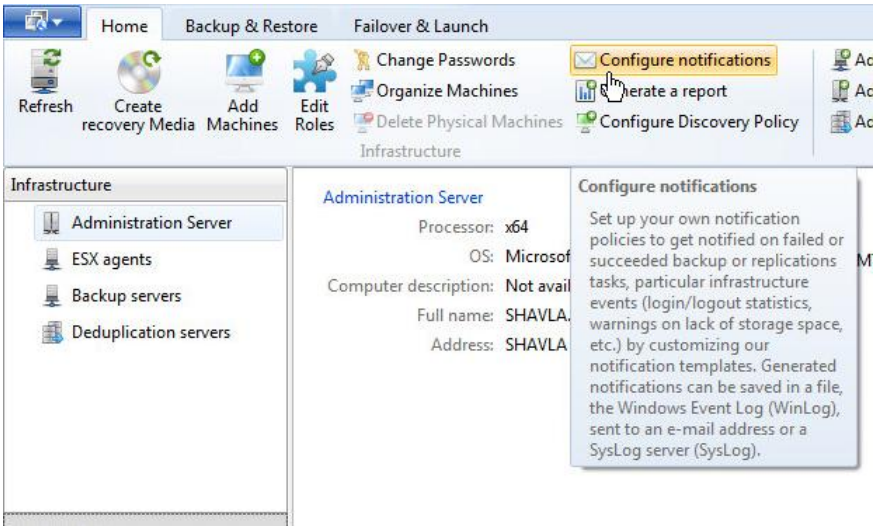

or go to **Notifications and Reports >** right click **Notifications**, then select **Configure notifications**.

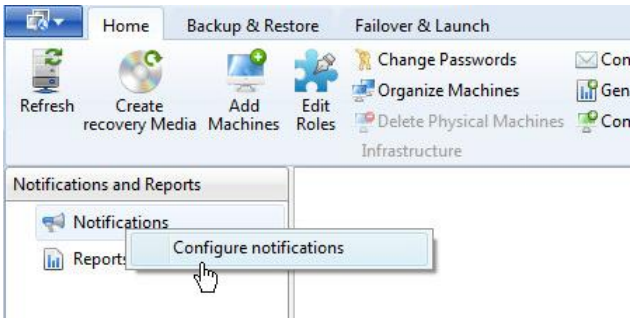

3. First you need to specify the required notification mode.

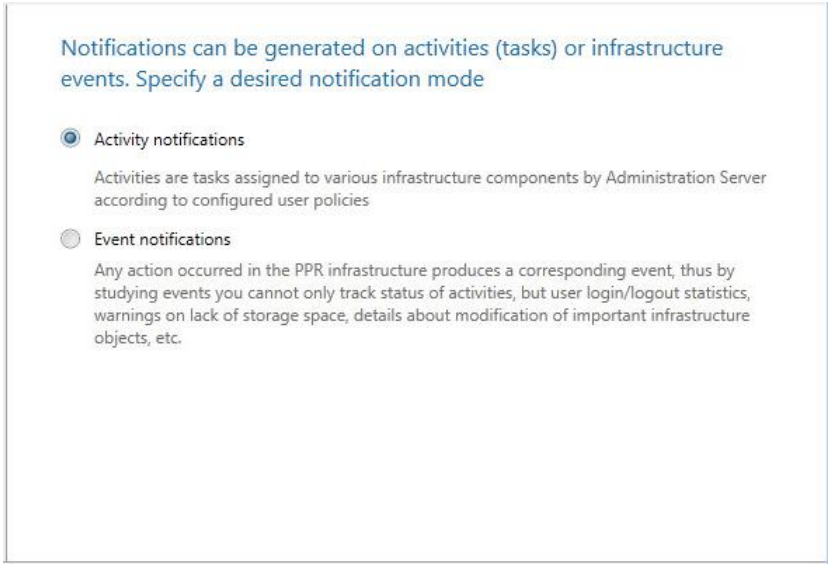

Infrastructure events you can be notified about are logically divided into two groups: **Activity notifications** and **Event notifications**. To choose the right notification mode, let's take a closer look at both of these options:

- Activities are tasks assigned to various infrastructure components by Administration Server according to configured user policies. In other words, after you create and launch a policy (backup, restore, retention, archiving, etc.), Administration Server determines and then assigns tasks to corresponding operation modules to perform your policy. Thus in the **Activity notifications** mode, you can set up notifications on tasks derived from all or specific user policies only.
- Any action occurred in the infrastructure produces a corresponding event. Thus in the **Event notifications** mode, you've got a lot more notification triggers to deal with beside tasks derived from user policies, e.g. login/logout events of PPR Console, warnings on lack of storage space, create/delete storage events, scheduler operation events, details about modification of important infrastructure objects, etc.

4. Next step the wizard prompts you to either use a ready-made notification template or create your own from scratch. Choose the first option if you're not advanced in PPR's architecture.

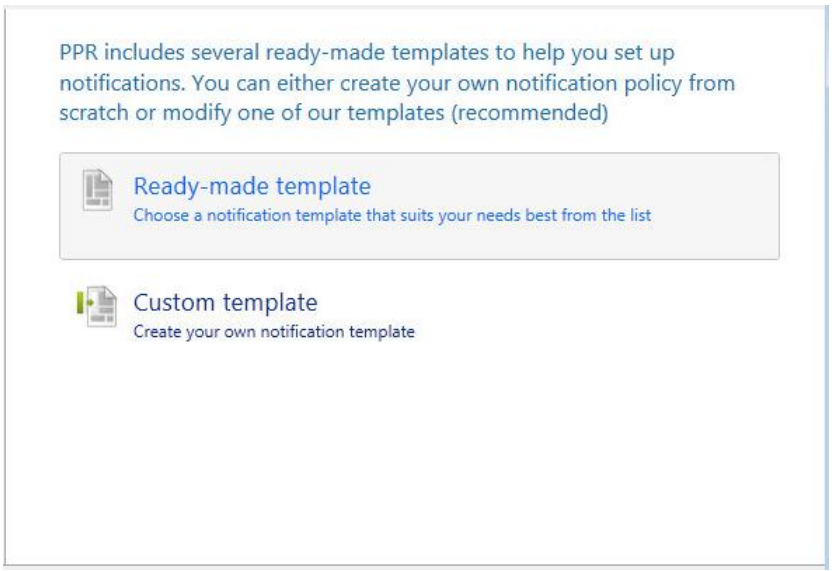

5. Choose the required notification template from the list. When you select a template, a small hint appears just below the list.

Activities templates:

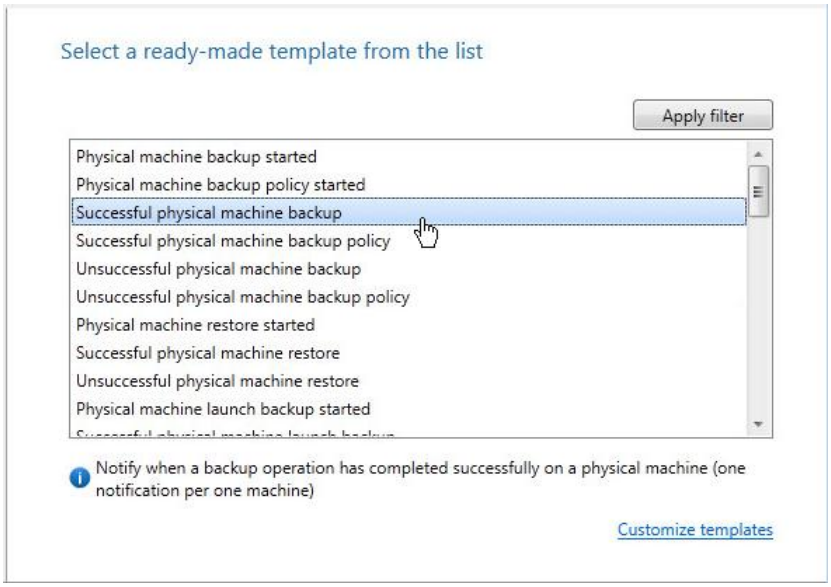

Events templates:

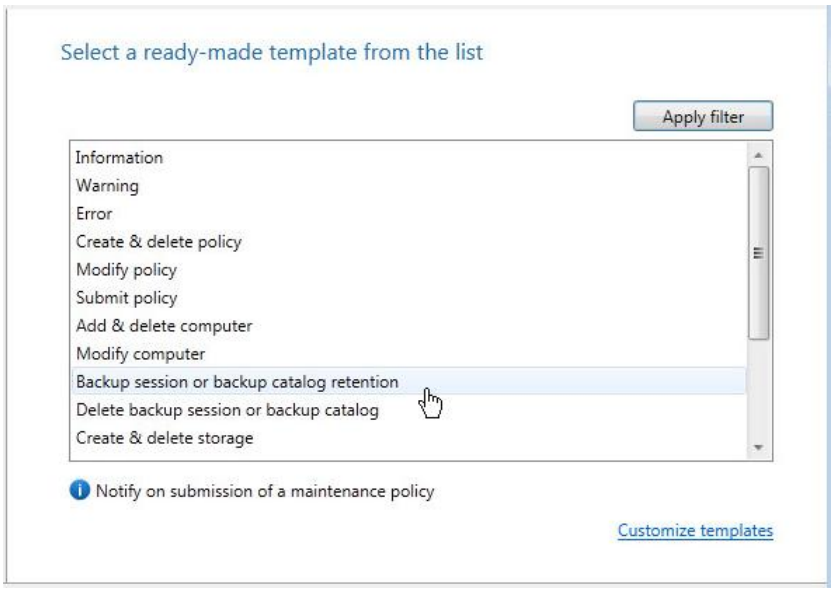

Use the **Apply filter** button to sort out one or several groups of templates you're interested in.

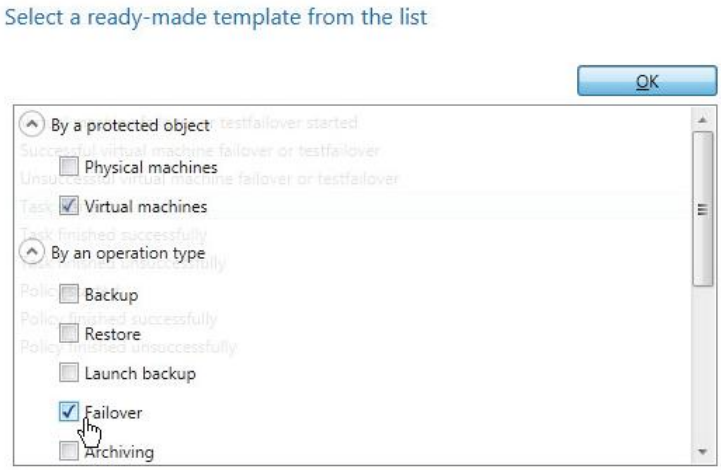

If you'd like to introduce some changes to the selected template, click on the **Customize templates** hyperlink.

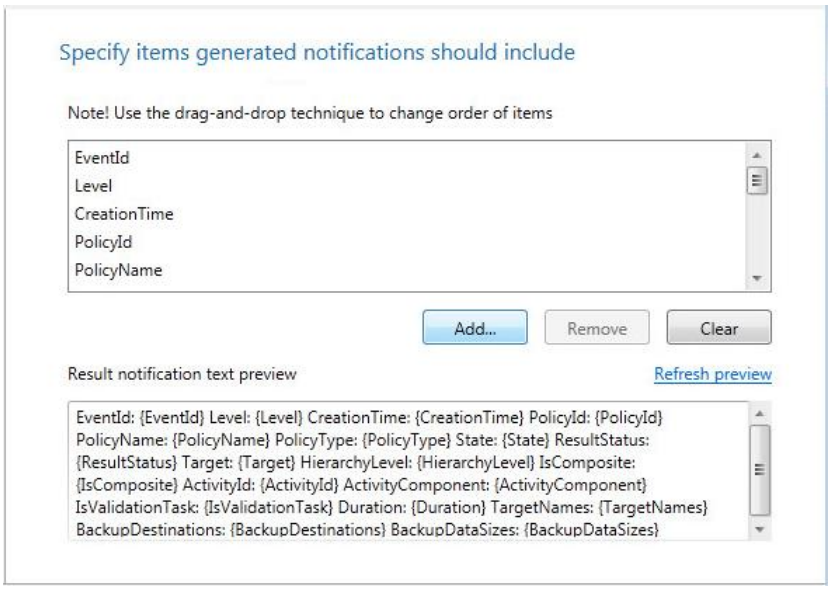

The opened dialog window is divided into two panes. The upper pane displays a list of notification triggers the selected template includes. Reorder triggers through drag-and-dropping to shape the notification body. When ready with amendments, click on the **Refresh preview** hyperlink to see the results in the lower pane.

If you're advanced in the PPR's architecture, you can optimize the template to suit your needs best by adding new or removing existing triggers. When you select a trigger, a small hint below the list may help you in the process.

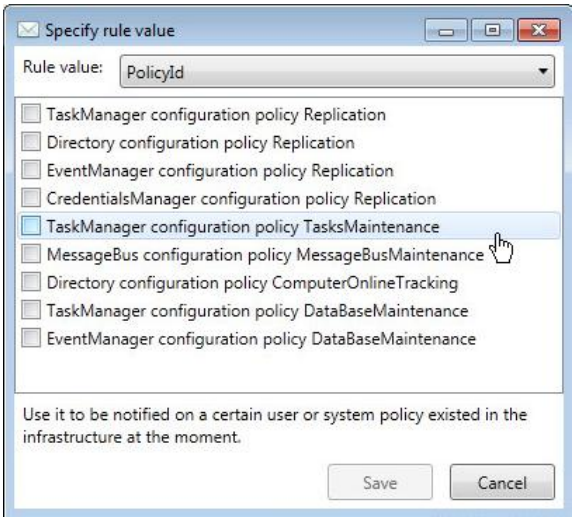

6. Select one or several pre-defined endpoints (collectors) where you would like the program to deliver generated notifications to. If you need to configure a notification collector at this stage, just use the corresponding hyperlink. For more information, please consult the Configuring Notifications chapter.

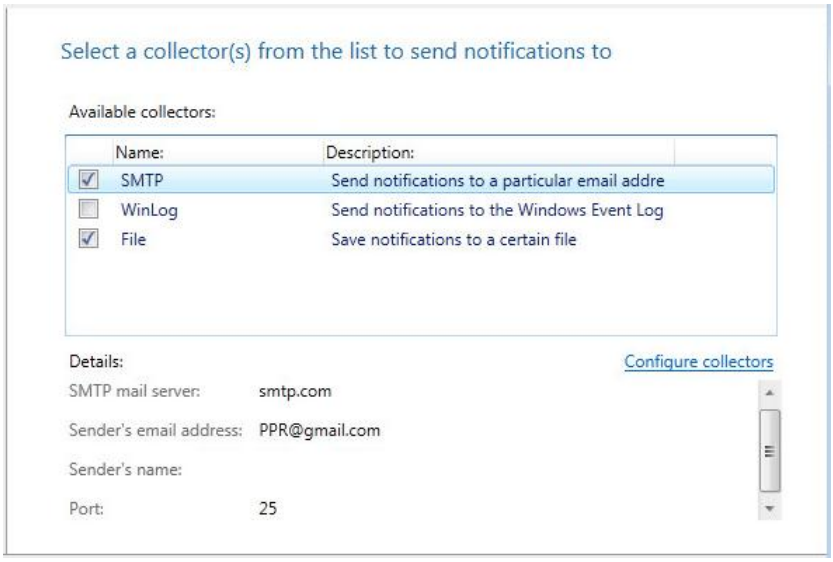

If choosing an SMTP-type collector, you will additionally need to specify an email recipient(s) and subject.

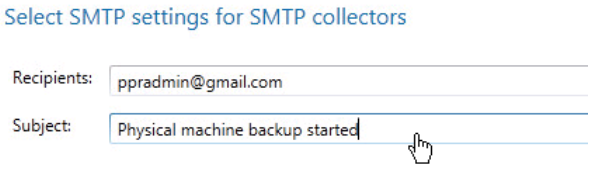

7. By default, the wizard offers to use the name and description of the selected template, which you can modify according to your requirements.

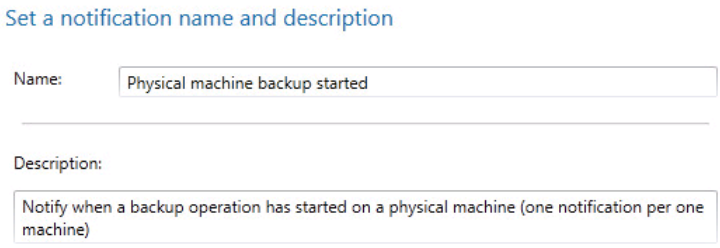

8. When you're done with all parameters click **Finish**. As a result you'll get a new notification policy. Use the corresponding options to disable, modify or delete it.

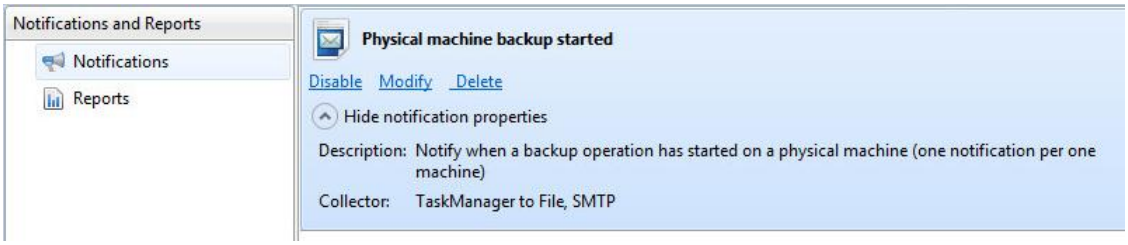

Reporting

Reporting is another tool of monitoring activities within the infrastructure from outside. You can set up automatic generation of reports by creating a corresponding policy (**Notifications and Reports > Reports**). As notifications, reports can be generated on activities of user created policies (backup, restore, storage archiving, backup data retention, etc.) and/or internal system policies (database replication, etc.). Reports can be generated on demand or by schedule (daily, weekly, monthly or for the last 7, 30, 90, 365 days) for succeeded or failed policies or both. In the current version of PPR, generated reports can either be emailed or saved to a file. There's only one available output format – html.

To set up automatic generation of a report, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Home** ribbon, then select **Generate a report**.

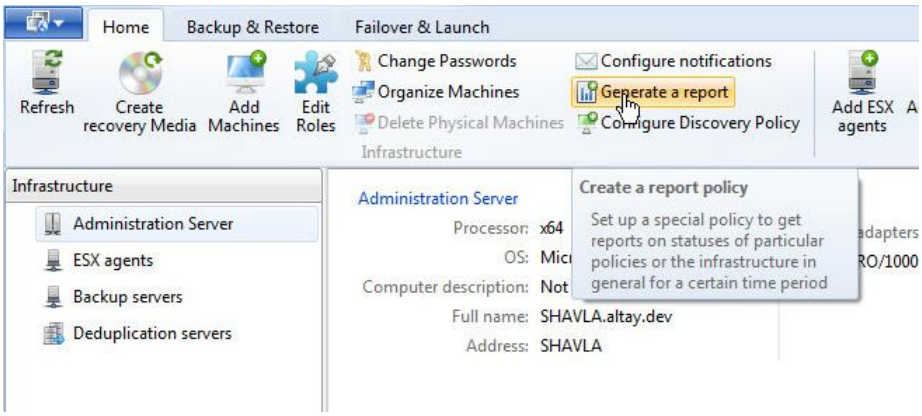

3. In the opened wizard choose a template for future reports. There are two available templates (**By policy** and **By policy type**). Let's take a closer look at each option to understand the differences.

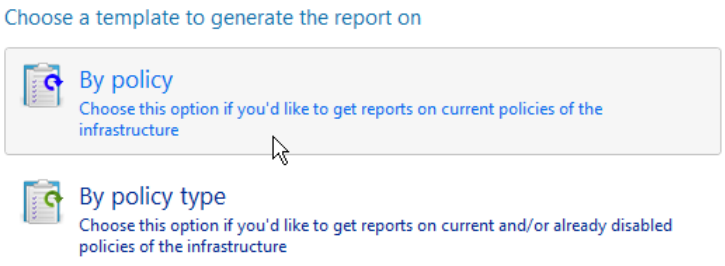

By policy:

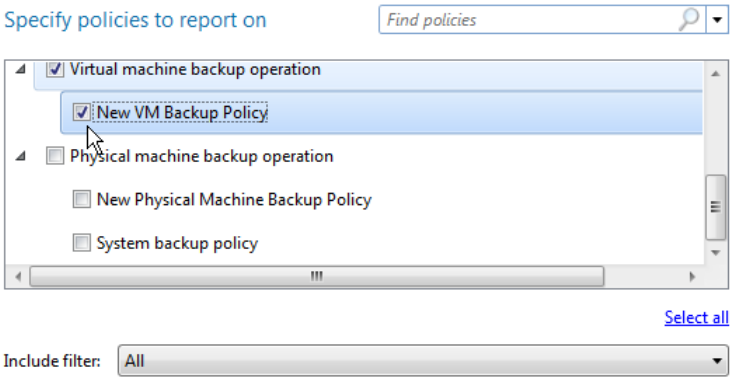

If picking the first option, the wizard will list all policies enabled in the infrastructure at the moment. Select those you need to be reported on, and then choose whether to report on failed, succeeded, or both types of policy runs in the **Include filter** option.

By policy type:

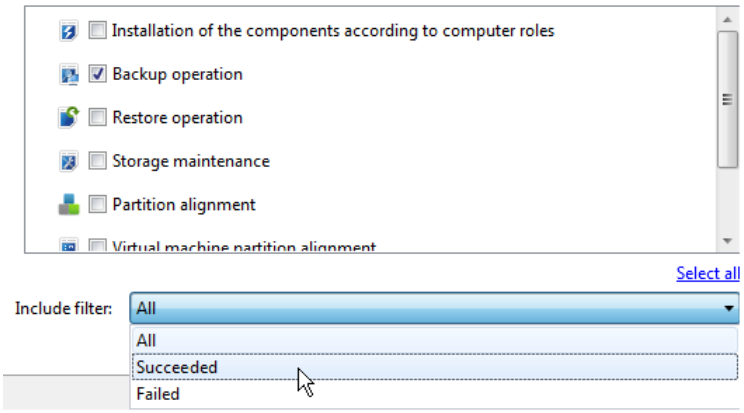

If picking the second option, the wizard will list all types of policies ever existed in the infrastructure (enabled and already disabled and deleted). Unlike the first option, you cannot pick a particular policy, but only a group of similar policies. Select those types of policies you're interested in, and then choose whether to report on failed, succeeded, or both types of policy runs in the **Include filter** option.

4. Next step you need to select the required report method, either **On demand** or **On schedule**. Let's again take a closer look at each option to understand the differences.

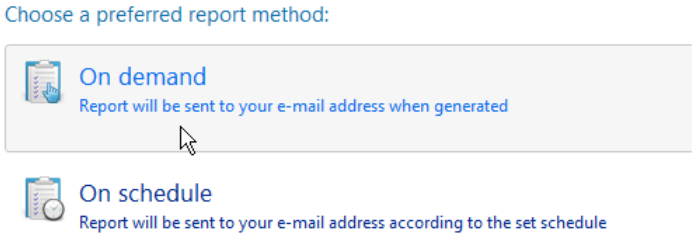

On demand:

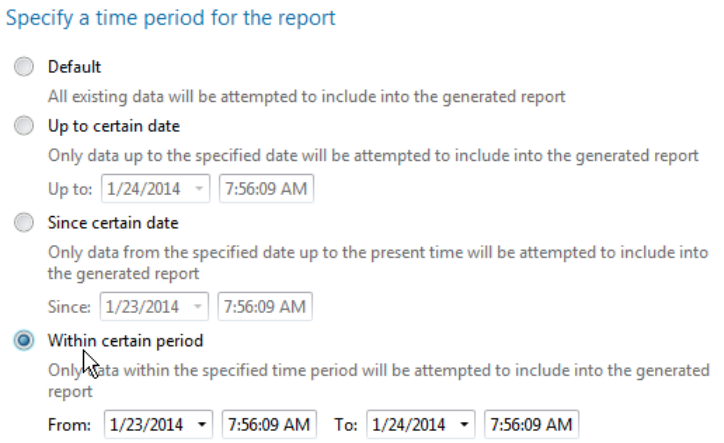

This option enables to specify a particular time period you'd like to be reported on right now.

On schedule:

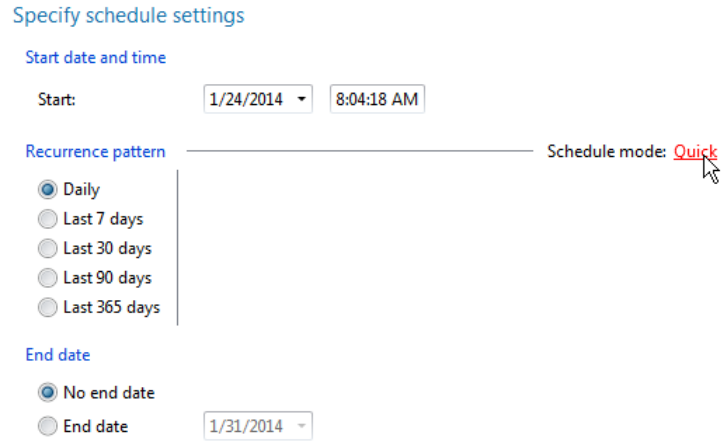

Select this option if you're interested in getting reports on a regular basis. There are two timetable modes available, i.e. **Quick** (for the last day, 7/30/90/365 days) and **Standard** (daily, weekly, monthly).

5. Now you've got to choose where reports should be delivered, either to a local file, or an email address. While the first option is obvious to deal with, let's consider how to configure reporting by email.

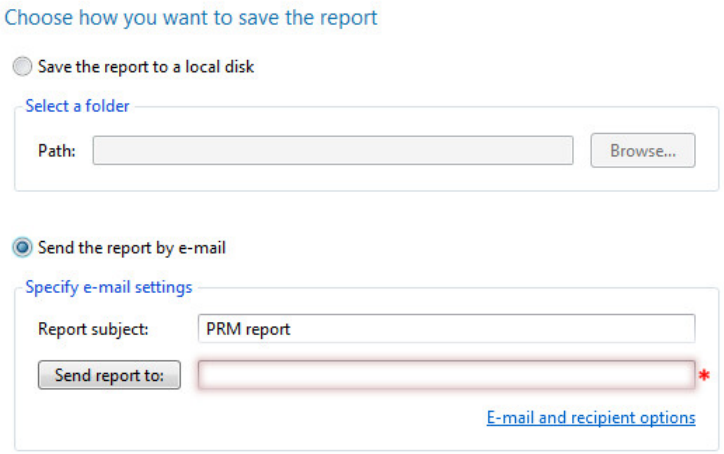

Give a subject to future reports, then either enter the target email address in the corresponding field manually or click on the **Send report to:** button to select an address from the list of previously used addresses. If the email

transport system has not been configured yet, click the **Email and recipient options** hyperlink to do it at this stage.

Email delivery (the first tab)

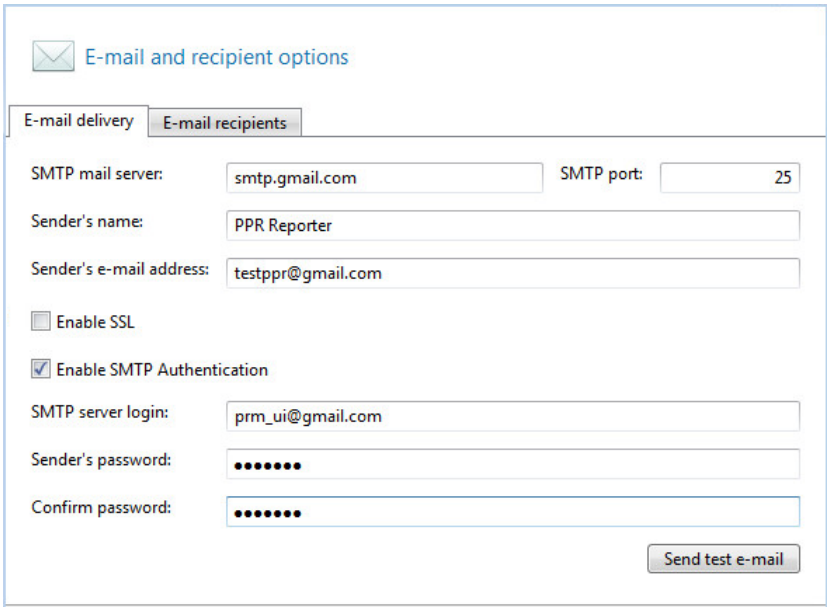

- **SMTP mail server**. To send notifications, it is necessary to have access to a computer running an SMTP (Simple Mail Transfer Protocol) server. All outgoing messages are first sent to the SMTP server, which in its turn delivers them to the required recipients. The address may be represented as a traditional Internet host name (e.g.: mail.com) or as an IP numeric address (e.g. xxx.xxx.xxx.xx).
- **Sender's name**. Provide a name of the sender.
- **Sender's email address**. Provide an email of the sender.
- **Enable SSL** (Secure Socket Layer). Activate the option to establish a secure connection to the email server.
- **Enable SMTP Authentication**. Activate the option to allow the program to authenticate on the SMTP server before sending messages, then provide valid user credentials in the corresponding fields.

Email recipients (the second tab)

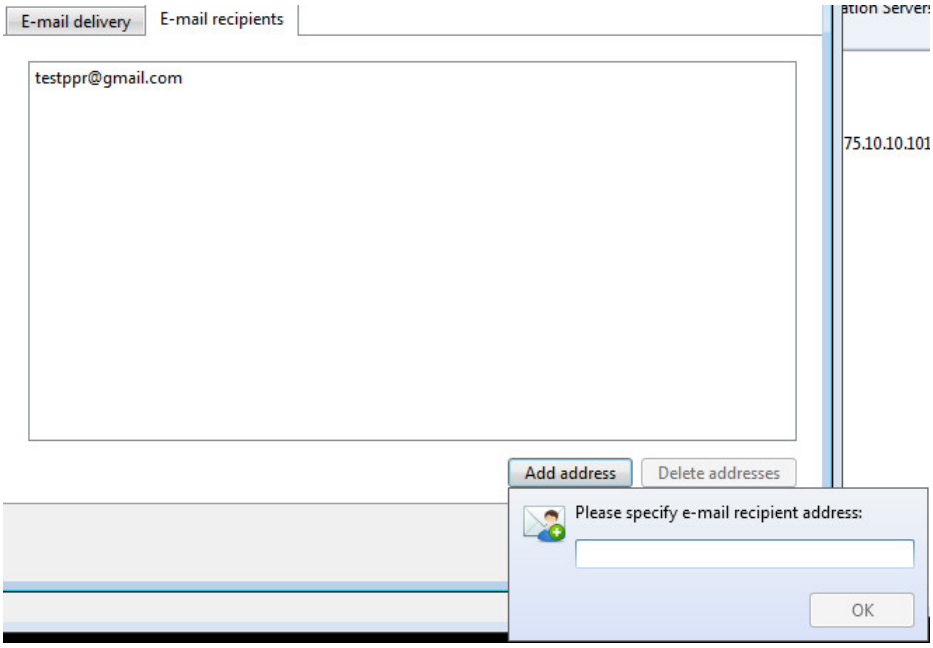

Use the corresponding buttons to add or delete recipients of notifications.

When you're ready with the settings, click on the 'Send test email' button to check if everything is OK.

6. When you're done with all parameters click **Finish**. As a result you'll get a new report policy. Use the corresponding options to force launch, view details on, disable, modify or delete it.

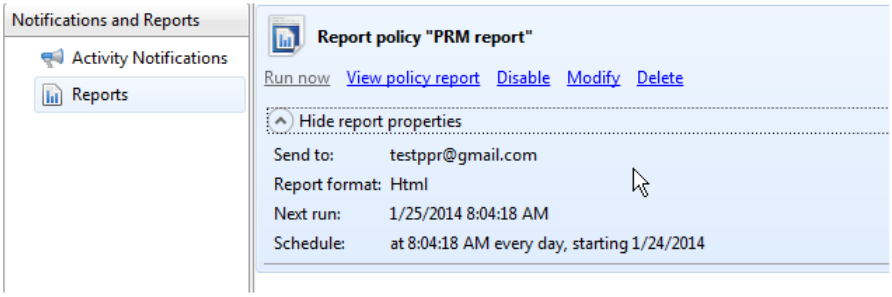

An example report:

PRM Task report

Report data from 2013-10-26 05:03 to 2014-01-24 04:03

Contents of report:

1. Policy: New Physical Machine Backup Policy
2. Policy: New VM Backup Policy
3. Policy: New VM Backup Policy
4. Policy: New VM Backup Policy (1)
5. Policy: System backup policy

Policy: New Physical Machine Backup Policy

Summary: Failed 0, Succeed 1, Total 1

StartTime	FinishTime	Duration	OperationType	Status	Results

Policy: New VM Backup Policy

Summary: Failed 0, Succeed 4, Total 4

StartTime	FinishTime	Duration	OperationType	Status	Results
12/24/2013 12:41:22 AM	12/24/2013 12:50:28 AM	9 minutes, 6 seconds	Backup	Finished / Success	- Size 0 Bytes, Image size 0 Bytes Backup type: Unknown: Success
12/24/2013 1:17:34 AM	12/24/2013 1:19:33 AM	1 minute, 59 seconds	Backup	Finished / Success	- Size 0 Bytes, Image size 0 Bytes

Updating the Infrastructure

The update procedure of the PPR infrastructure is simplified as possible:

1. Download a new version of the product from Paragon's website, the **My Account** section.
2. Launch the obtained installation package on a machine where **Protect & Restore Server** is deployed.
3. Read and accept all conditions of Paragon's license agreement by selecting the appropriate option.
4. Select **Upgrade**.

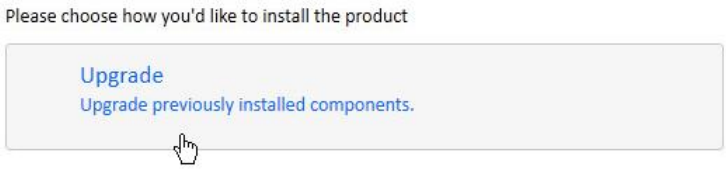

5. Provide credentials of a domain or a local administrator. Please make sure the domain administrator also joins the 'local admins' group. Click **Install** to initiate the upgrade process.

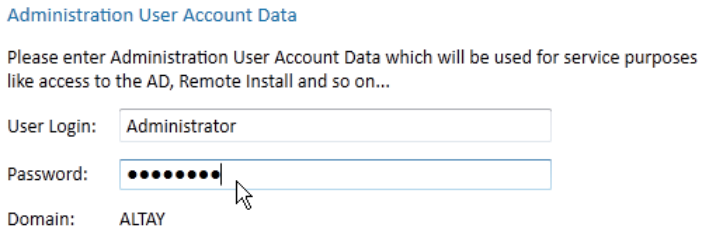

Beside domain or local administrators, PPR can also be administered by members of special groups created during the installation. For more information, please click here.

6. You will be warned if the currently installed version of PPR doesn't support rollback to the previous state. Click **Yes** to confirm the operation.

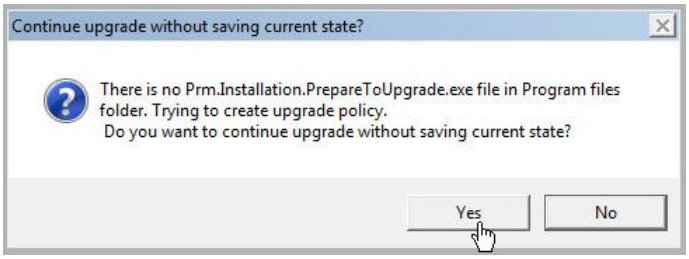

7. The wizard will copy updated MSI packages to the repository of Installation Server and update versions of PPR Server and PPR Console (if installed on the same machine). When the update process is over, click **Finish** to exit.
8. Launch PPR Console. If the console is installed on a different machine from PPR Server, please use the PPR Installer to update it, otherwise you won't be able to log in.
9. If a connection with the server has been established, you will be notified through a popup window that a new product version is available, and the corresponding icon on the **Home** ribbon will become active. Click **Update > Update System** to initiate update of other infrastructure members.

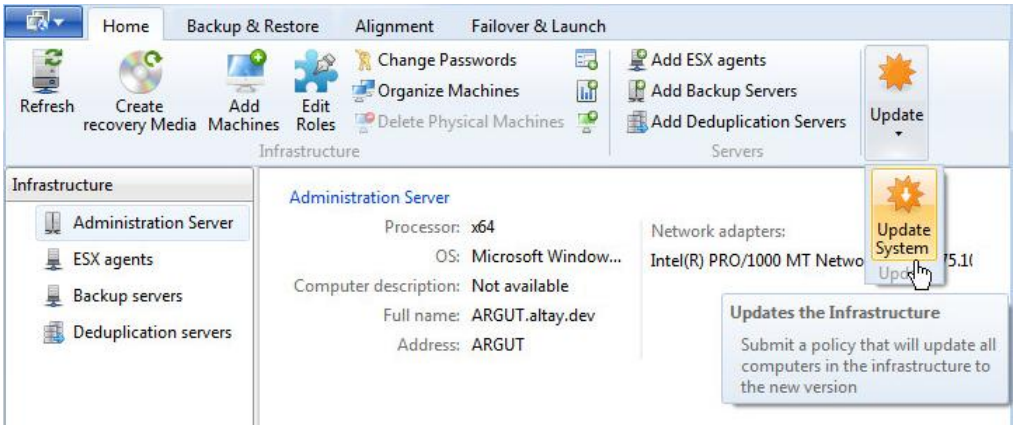

Replication of infrastructure databases is accomplished every 5 minutes, thus information on the new version will get to the console during this period.

10. Choose whether to accomplish the operation immediately or set a timetable for later execution.

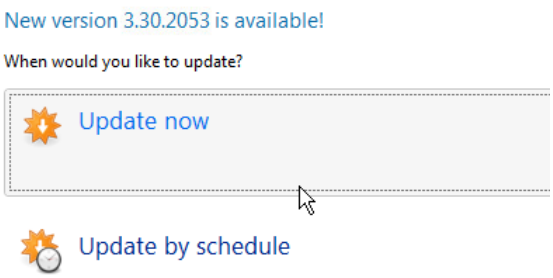

11. You will be notified on the oncoming actions. Click **Update** to confirm the operation.

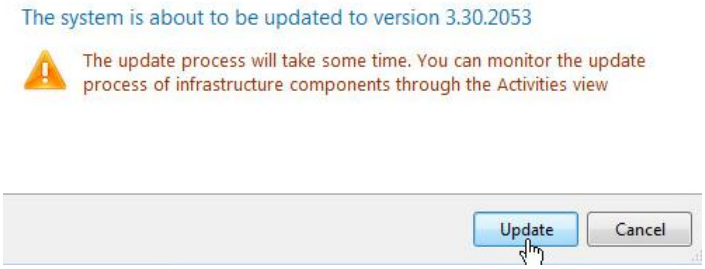

12. Open the **Activities** pane to monitor update of other infrastructure members.

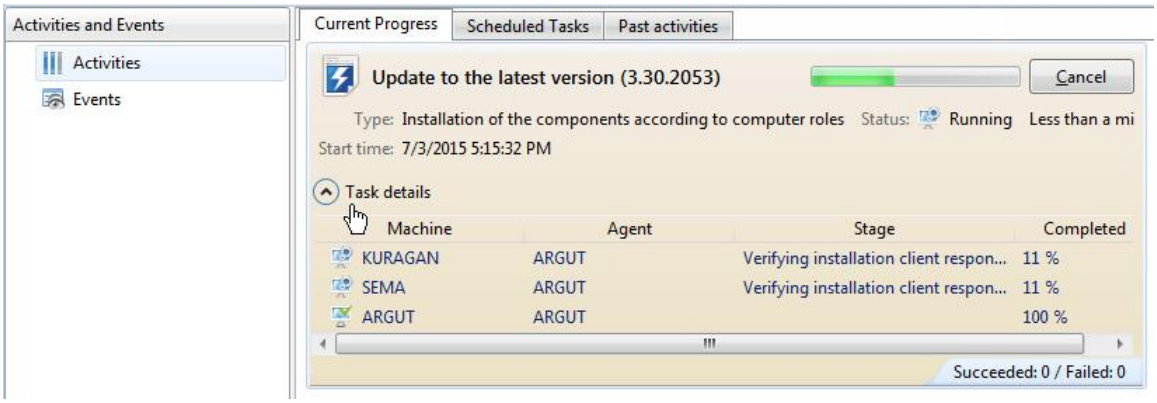

13. When done, check versions of all PPR infrastructure members by selecting **Machines > Physical Machines**.

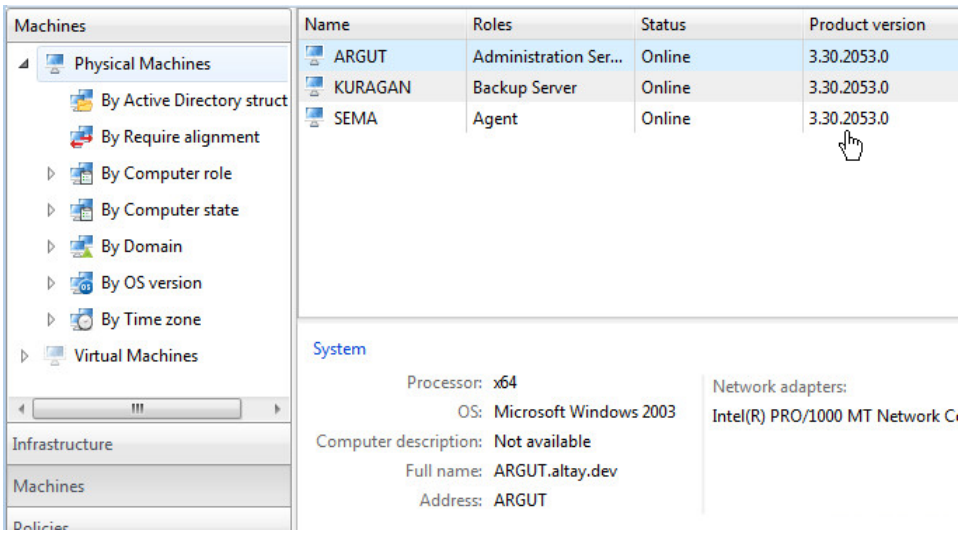

Enabling Role-based PPR Security

If you decide not to use the embedded security facilities when deploying PPR, machines that join the infrastructure and all traffic get vulnerable to unauthorized access. We consider protection of ten and more production servers and workstations in the non-security mode unacceptable. To switch on the Role-based PPR security for all infrastructure members, please do the following:

9. Launch the installation package on a machine where **Protect & Restore Server** is deployed.
10. Select **Change Security**.

11. Select whether MS Active Directory facilities (**In Domain**) or credentials of a local machine (**In Workgroup**) will be used to authenticate users to grant access to the PPR infrastructure.

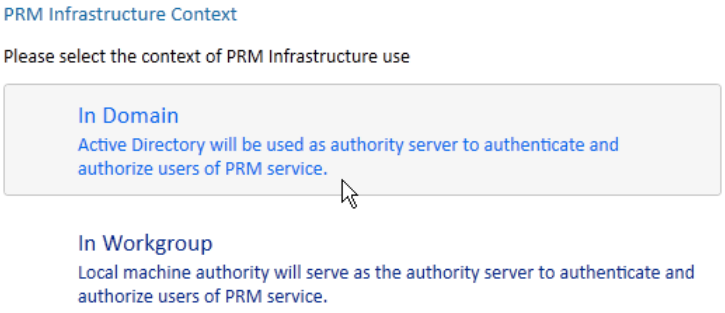

12. Depending on your choice at the previous step, provide credentials of a domain or a local administrator. Please make sure the domain administrator also joins the 'local admins' group.

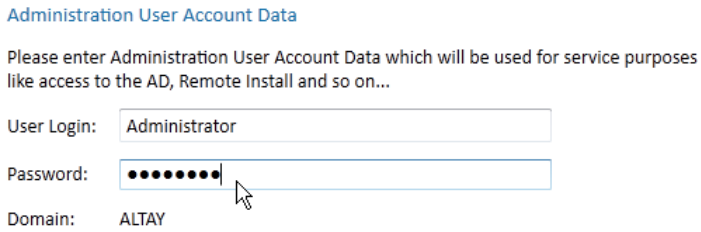

Beside domain or local administrators, PPR can also be administered by members of special groups created during the installation. For more information, please click here.

13. If selecting Active Directory environment (**In Domain** option), the setup wizard will display all of its organizational units (OU) prompting you to choose where PPR's organizational unit with the default security groups should be created.

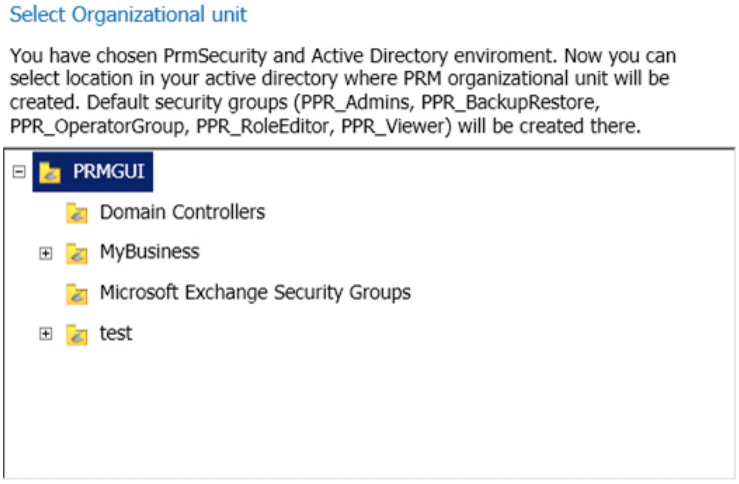

14. Once confirmed, the wizard will first assign the "RolePending" state to all machines of the infrastructure to avoid faulty exchanges between machines while the role-based security is being enabled. Then it will restart the PRM service to reconfigure it according to the selected security mode. When done, there will be created an installation policy called "Inversing security mode", which you can see by going to **Activities and Events > Events**. This policy will reconfigure the PRM service on all other machines of the infrastructure.

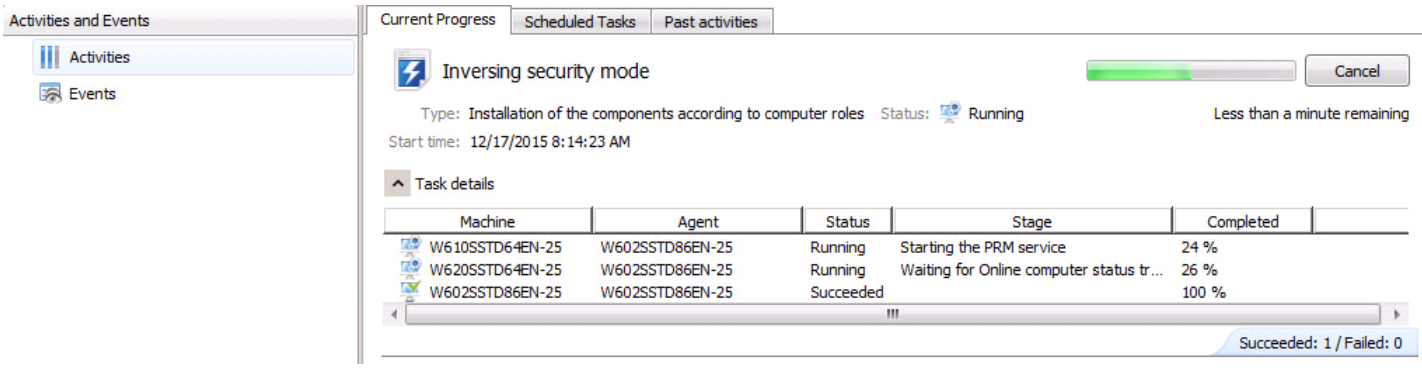

Removing Machines from the Infrastructure

1. Launch Protect & Restore Console.

2. Select **Machines > Physical Machines**, then a computer (use “Ctrl” or “Shift” for selecting several machines) you’d like to be removed from the infrastructure. Use the **Delete Physical Machines** icon to initiate the operation.

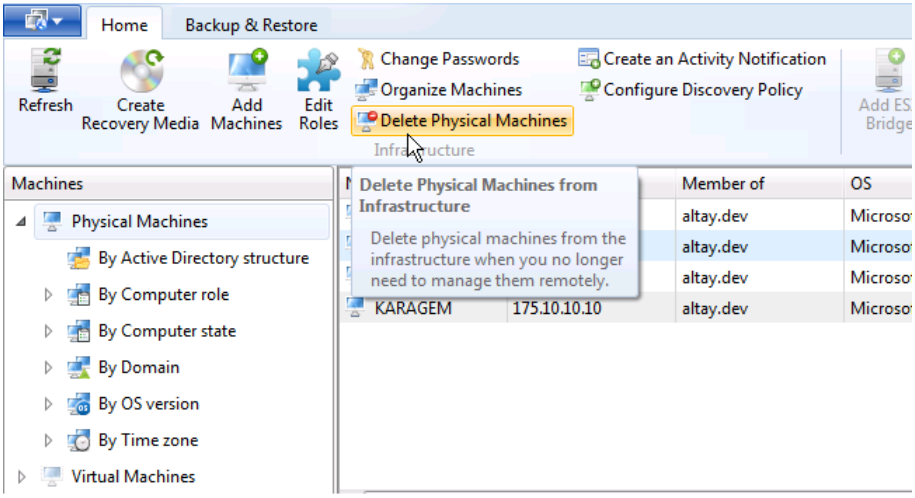

You can also initiate this operation by the right click of the mouse button on the required infrastructure member, then selecting the corresponding option. In this case you won’t need to specify a machine you’d like to work with.

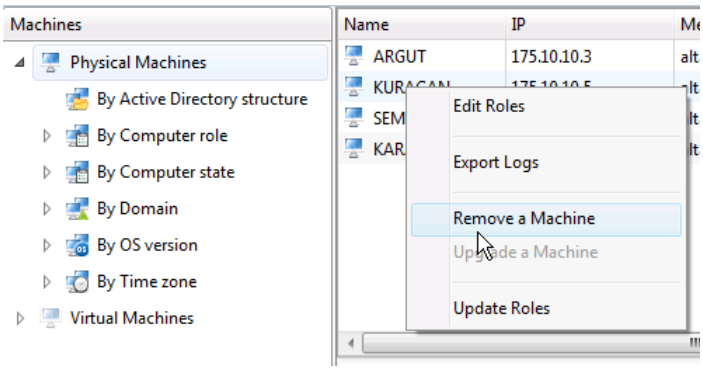

3. Confirm the operation in the opened dialog. Please note that this operation initiates de-installation of all components of Protect & Restore on the selected computer.

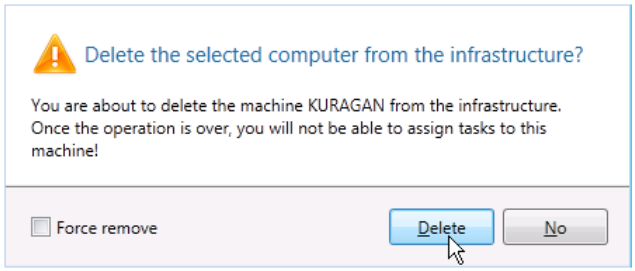

4. To get detailed information on the operation progress, click on the link in the popup window or select **Activities and Events > Activities**, or **Activities and Events > Events**, where you can see all events of the infrastructure. Just click on the interested one to see details. Use the **Refresh** icon to force update of the information. To know more on the subject, please consult the corresponding chapters: Managing Activities, Managing Events.

If the delete operation fails for the selected target machine, please reopen the dialog and additionally mark the “Force Remove” option to delete all information on this computer from the infrastructure. Then uninstall all components of PPR on-site by using the Installer

and its corresponding section.

Collecting Logs

In case of having difficulties with handling the product you can address our support engineers for assistance. To submit a support ticket, first you need to collect all operation logs. PPR simplifies this operation with a handy dialog.

To prepare a package of operation logs, please do the following:

1. Launch Protect & Restore Console.
2. If a connection with the server has been established, click on the **Logo** button, then select **Export logs**.

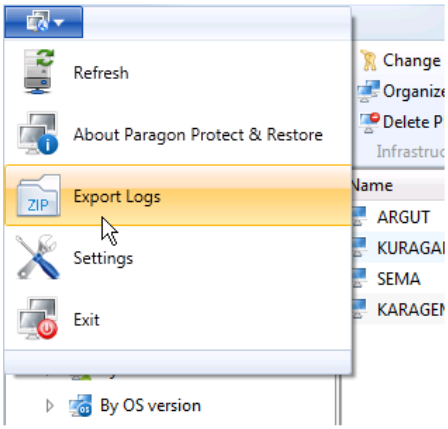

3. If you know when the encountered problem first started, please specify a time period to collect logs for. That will help you to minimize the created logs package. If you're not sure, leave the default option as is. Please note that the entire logs database may be several gigabytes in size.

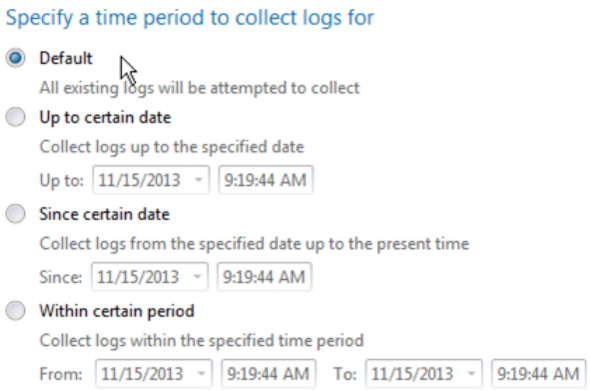

4. Select infrastructure members, which logs you're interested in. Please note logs from Administration Server can help to figure out and resolve 90% of issues. However when the database replication with one or several infrastructure components is not possible, or has not been completed yet, the logs stored on Administration Server will not contain the most up-to-date information. Click **Next** to proceed.

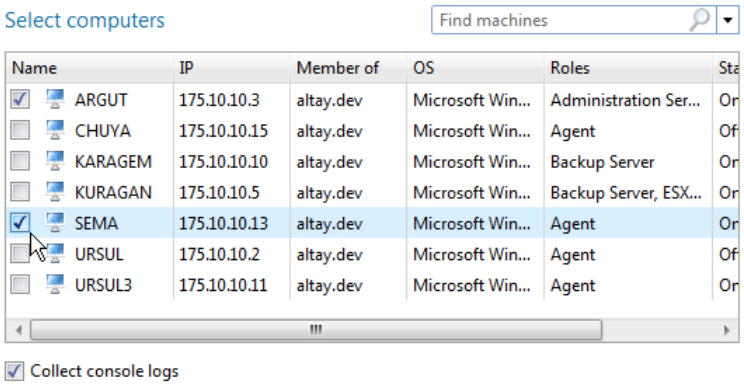

5. Browse to a place where you'd like to save the logs to. If you'd like Windows Explorer to open in the specified folder once the operation is over, please additionally mark the corresponding option.

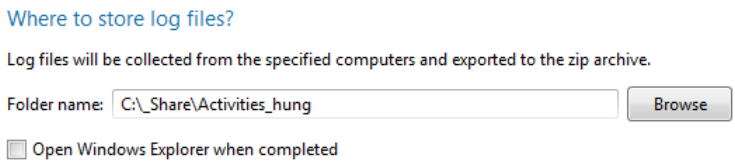

6. Choose whether to collect logs on activities of all PPR users or only those you're interested in. Click **Export** to initiate the operation.

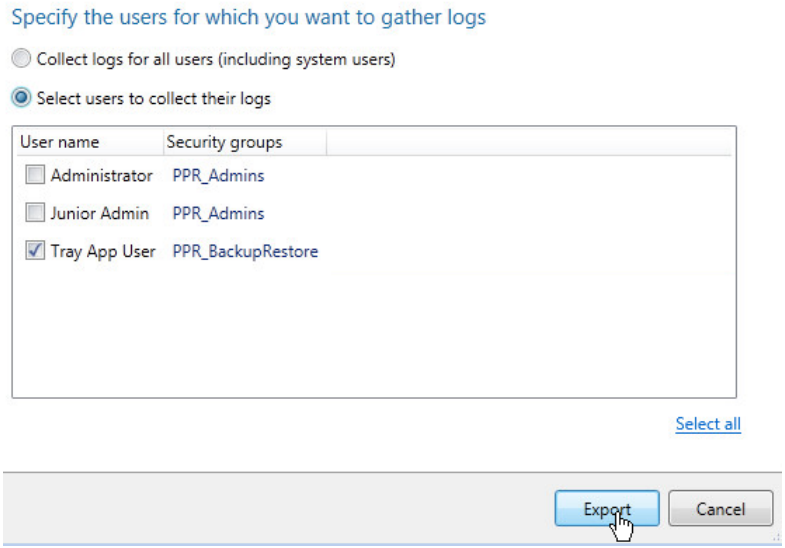

Log files do not contain any confidential information on the operating system settings or the user documents.

Known Issues

1. There's information in this manual that the user can manually install other infrastructure components besides Protect & Restore Server, Console, and Backup Agent. But I don't find corresponding options in the PPR installer to do that.

You're quite right. The current version of the PPR installer doesn't allow installing other components, but PPR Server, Console, and Backup Agent. Please install all other components remotely from PPR Console.

2. Unable to add a USB 2.0 flash stick to a disk pool with the following notification: '**Adding device with multiple volumes is not allowed…'**

   Most probably your flash stick has been formatted as "superfloppy", when entire removable media is treated as a single partition and there's neither GPT nor MBR partitioning scheme. It has a special signature, which has not yet been supported by PPR. Please re-format the problem flash stick with DiskPart to fix the issue.

3. My network backup storage fails regularly with the following error messages:

   'Insufficient system resources exist to complete the requested service…'

   'The specified network name is no longer available' and additionally **'The network path was not found'**

   Most probably your NAS has limited capabilities and the maximum number of file connections for each public shared folder has been restricted. PPR requires more than the maximum allowed connections during its work, which leads to the above problem. Please set this value on NAS to 0 to fix the issue.

4. A physical machine restore policy completes with one of the following errors: '**Failed**' or '**Finished, there were errors**'. When studying corresponding events, you find this exception: **'One of the target disks is smaller, than that of the source'**

   Most probably the target disk(s) is smaller in capacity than the source. PPR cannot yet restore with shrink, please use disks of the same capacity or larger than the source.

5. A physical machine restore policy completes with one of the following errors: '**Failed**' or '**Finished, there were errors**'. When studying corresponding events, you find this exception: **'Cannot build and merge replica…'**

   If the target machine resides in a different subnet, when you start it up from our WinPE media, it acquires a temporary name 'minint-…', which isn't registered in DNS of your Active Directory, thus Administration Server and Backup Server can't resolve this name, since a broadcast packet with wins name has been blocked. To fix the issue, please register 'minint…' in your DNS manually to accomplish restore.

6. PPR fails to do backups or replicas of virtual machines hosted by VMware vSphere with the following error: **'VIM SDK plugin is not loaded'**

   Most probably you've deployed ESX Agent/Backup Server on a machine running a 32-bit Windows OS. Please use a 64-bit Windows OS machine, as PPR only includes 64-bit ESX Agent. If you don't use ESX-based storage, then you're free to deploy Backup Server on a 32-bit machine.

7. I want to keep little backup data on primary storage (5 sessions only), while much more data on secondary storage (4 weeks). Therefore I specify data retention for my backup policy to 5 sessions and 4 weeks – for the entire secondary storage. However, I notice that my secondary storage contains 5 latest backup sessions only. What do I do wrong?

   In PPR backup policy retention settings override general backup storage settings. Besides, when backup data is being moved from primary to secondary storage, backup policy settings (including data retention options) are being transferred 1:1 as well. That's the cause of your troubles. To fix the issue, please do not use backup policy retention settings, but set short lifetime (5 backup sessions or whatever you need) for the entire primary storage.

8. I've chosen the 'AD Security' mode, but it doesn't work properly.

   The 'AD Security' mode is designed for pure Active Directory domain environments. If you've got at least one non-domain machine (a work group machine based on Windows or Linux) you're going to protect or deploy one of PPR's components on, please use the 'PRM Security' mode. Click here for more information.

9. I've noticed that gradually some operations in the PPR GUI Console require much more time to accomplish.

   If your PPR infrastructure is working in the 'AD Security' mode, then either install the console on a domain machine that does not have the role of Administration Server, or mark the "Use Windows session credentials" option at the console startup.

10. After restoring a domain controller, all machines of the PPR infrastructure has gone offline.

    That's quite normal, if your PPR infrastructure is working in the 'AD Security' mode.

11. Restore through the WinPE environment fails with the following error: "Either the target name is incorrect or the server has rejected the client credentials". Access credentials are valid, checked it several times.

    Most likely the target machine has two network cards onboard. Please disable the NIC that has access to external network and try the restore operation again.

12. When trying to open an event attach I've got the following error: "access is denied".

    If your PPR infrastructure is working in the 'AD Security' mode, please run the PPR console on behalf of the domain administrator or a member of the 'PRM ADMINS' group with the privileges of the local administrator.

13. I cannot log in to the PPR Console under a local administrator that joins the 'PRM Admins' group.

    Most likely you're using the 'AD Security' mode. If this is your case, please log in to Windows OS under some other administrative account, and then try to open the console again.

14. After changing access credentials on my Windows domain controller, PPR fails to access remote infrastructure members.

    Most likely the PPR Console is installed together with Administration Server. To fix the issue, please install the console on a domain machine that does not have the role of Administration Server.

15. The current version of PPR allows attaching registered storages from another PPR infrastructure, which is no good, as it may cause conflicts.

    If you'd like to attach storage from another PPR infrastructure, please make sure it's unregistered (deleted) and won't be used by two PPR infrastructures simultaneously.

16. I cannot attach my ESX storage as the wizard finds no storages in the specified location.

    If ESX storages you're trying to attach have been created in a previous version of PPR (3.23.2289 or earlier), then you won't be able to accomplish this operation at all. You can see the product version by clicking the **Logo** button, then **About Paragon Protect & Restore**.

17. I've got problems attaching my ESX storage.

    Most likely you've moved your ESX storage to another datastore. This scenario is not supported at the moment.

Appendix

PPR and Windows Firewall

During installation of the following PPR components:

- PRM30_Common_ea_x32.msi (x64)
- PRM30_InstallationClient_ea_x32.msi (x64)
- PRM30_PAT_Console_ea_x32.msi (x64)

Exceptions are automatically added to Windows Firewall for the corresponding applications/services:

- Prm.Common.Server.exe
- Prm.Installation.Client.exe
- Prm.Console.Shell.exe

Each exception rule is configured with the following parameters:

- True for all network profiles (domain, personal, public),
- Allows the rule for any application and/or service,
- Allows the rule for any IP address,
- Allows the rule for any protocol,
- Allows the rule for any port.

Rules are added only if the **Windows Firewall/Internet Connection Sharing (ICS)** service is being started at the moment of the PPR components installation, no matter whether Windows Firewall is enabled or not. If the mentioned service had been stopped during the installation, and then started by the administrator, Protect & Restore Console would not be able to connect to Administration Server with the given error: '**Failed to connect to <Administration Server>. Please make sure the PRM service is running on the remote computer, or try again later.**' This actually means there has not been added an exception to Windows Firewall for **Prm.Common.Server.exe**. The only way out is to manually add exceptions rules for the mentioned above application and services as described in this section.

Glossary

DAG – Database Availability Group

A database availability group is a high availability and data recovery feature of Exchange Server 2010. It can consist of up to 16 Exchange mailbox servers, automates recovery at the database-level after a database, server or network failure.

DAGs replaced the high-availability model of Exchange Server 2007, which was based on local continuous replication (LCR), standby continuous replication (SCR), single copy clustering (SCC) and cluster continuous replication (CCR).

SCC - Single Copy Cluster

This is a minimum two node cluster that relies on Microsoft Failover Clustering Services. This requires shared disk like a SAN and has a single copy of the data.

LCR - Local Continuous Replication

This is a single Exchange Server 2007 solution to provide data redundancy. You can run all of the server roles on this however there are some caveats for public folders. To be effective this solution requires two external drive arrays and two array controllers to provide true redundancy.

CCR - Cluster Continuous Replication

This is a two node cluster that relies on Microsoft Failover Clustering Services on Exchange Server 2007. This does not require shared disk however would require a "witness" node. Other than the mailbox role, no other roles can be installed on the cluster. There active node of the cluster replicates all changes to a passive copy of the database. A minimum of three Exchange servers would be required (2 mailbox nodes and a non-redundant Client Access and Hub Transport).

SCR - Standby Continuous Replication (SP1)

This allows replication of databases to other Exchange Server 2007 located anywhere on the Intranet. This replication can be done to and from any type of mailbox node (other than a mailbox server that uses LCR).

SIOS - Single Instance Object Storage

It’s a method of redundancy elimination, when, before running a full backup, it is analyzed if any identical data block already exists on the backup storage to process and store one copy of any block only, thus minimizing the backup storage requirements.

High Availability and Site Resilience

It’s a new replication technology implemented in Exchange Server 2010 (see http://technet.microsoft.com/en-us/library/dd335211.aspx).

VSS - Volume Shadow Copy Service

It’s a technology that provides the copy/backup infrastructure for the Microsoft Windows XP/Vista/7/Server 2003/2008 operating systems. It offers a reliable mechanism to create consistent point-in-time copies of data known as shadow copies. Developed by Microsoft in close cooperation with the leading copy/backup solution vendors on the market, it is based on a snapshot technology concept.