Configuring Backup Storages in Paragon Protect & Restore

Best Practices

last updated: August 2013

Overview

Paragon Software's Protect & Restore (PPR) offers a unified system and data protection solution for virtual and physical machines. PPR delivers comprehensive agentless protection for virtual environments hosted by VMware vSphere or standalone ESX servers[1], and agent-based protection for physical and virtual Windows systems on any hypervisor. As its backbone, Paragon leverages a patent-pending distributed architecture allowing for efficient centralized and remote management of hundreds or even thousands of machines on the network.

In this document we will touch upon Paragon's approach to the problem of dual protection, highlight the benefits of storage replication or archiving over competitive parallel backup technologies, consider various types of storages PPR supports and the ways how contents of storages can be imported from one PPR infrastructure to another.

Paragon's Two-tier Storage Infrastructure

PPR supports a two-tier storage infrastructure that provides for minimization of backup windows and network traffic for simultaneously made backups. In this type of infrastructure, the PPR administrator 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, SFTP, 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.

When compared to competitive pure parallel backup technologies, Paragon's approach is obviously more efficient as it enables to:

- Avoid extra snapshots of ESX guests or physical Windows machines. Thus it minimizes backup windows, releasing CPU and network bandwidth resources.
- Schedule transfer of backup data to secondary storages when most appropriate.
- Do on-the-fly conversion of backups to replicas or vice versa depending on the target secondary storage type.
- Configure archiving from one secondary storage to another, thus opening up a path to a three, four,…-tier storage infrastructure.

In the GUI console, the PPR administrator should take the following actions to set up archiving of backup data to secondary storage:

[1] Agentless Hyper-V support is scheduled for December 2013.

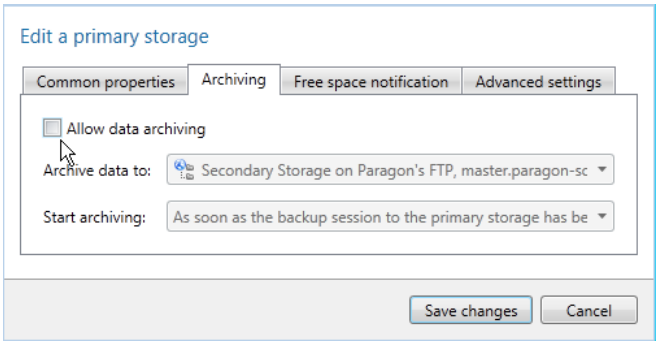

- Call properties for a primary or secondary storage, which backup data needs to be archived;
- Allow the data archiving by marking the corresponding option;
- Specify the target secondary storage, where backup data should be archived to;
- Set up an operation timetable. By default, archiving will be triggered once a backup session on the source storage has been completed.

Supported Storage Types

As mentioned above PPR differentiates primary and secondary storages, each having its own purpose. Primary (first-tier) storage is assumed to reside as closer to target machines as possible to ensure the highest backup or replication performance – it can either be a local folder of Backup Server or a network share for backup images of virtual and physical machines or ESX datastore for VM replicas.

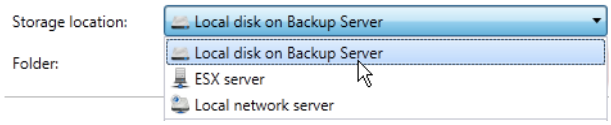

Secondary (second-tier) storage is assumed to be offsite. Paragon's approach implies no direct use of this type of storages, but only through special archiving policies. PPR enables to reside a secondary storage in a local folder of Backup Server, ESX datastore, a network share, or an FTP/SFTP server.

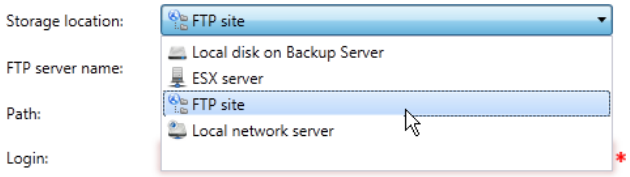

Beside general properties common for all types of storages (name, retention schedule, etc.), there are unique properties, relevant for particular type of storages only (credentials, pool and datastore of ESX server, FTP server name, port and credentials, a network share, etc.).

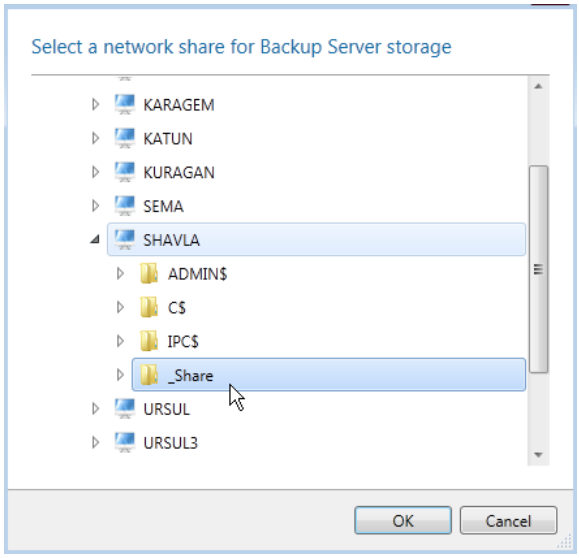

Storage Maintenance

PPR includes a number of tools to maintain storages and backup data they contain:

- Almost any property of an existing primary or secondary storage can be reconfigured at any time. Among other parameters, 'Retention policy' and 'Retention schedule' are those that require a bit of explanation. A **retention policy** specifies how long backups (replicas) should be kept on the storage. By default, for every target machine PPR creates a full backup (replica) during the first run, then incremental updates according to corresponding backup policy schedules. When time comes, all restore points beyond the set limit are merged with their full backup, thus creating a new full backup. A **retention schedule** specifies how often backup data inside the storage will be checked for and processed according to the specified data retention policy (every day, by default). Thus both these options enable to automate removal of obsolete backup data on a particular storage to free up space for new backup data.

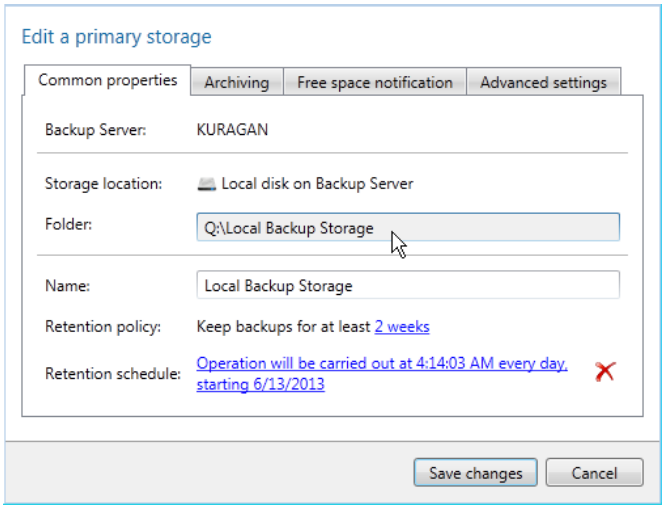

- For easier administration of backup data, all storages are open for browsing. Having a list of all created backup items at hand, the PPR administrator can easily find and initiate restore of required backup items, run an integrity checkup for those that are critical, or delete those that are not needed any more. When deleting an increment from somewhere in the middle of the incremental chain, PPR automatically initiates data merging to keep the whole chain consistent.

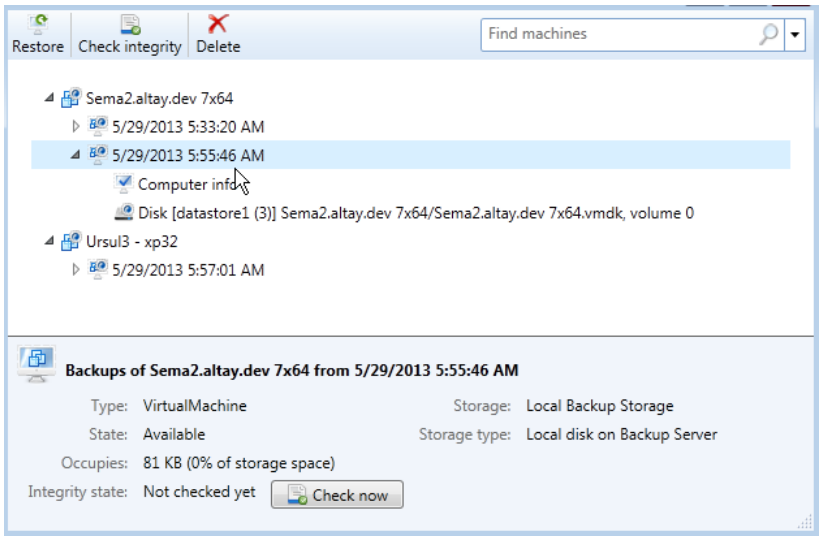

- Deleted storages (by default, backup data is not deleted when removing a storage form the infrastructure) or storages from another PPR infrastructure can be easily imported (attached) to the existing infrastructure to use backup data they contain. Most parameters are identical to creation of a new storage, except for the PPR administrator should provide the required storage role (primary or secondary) and a local or network path to the existing storage, which backup data is to be imported.

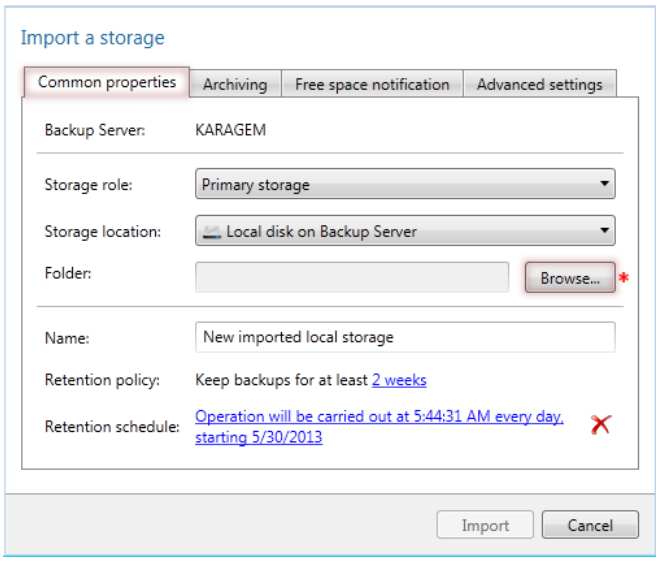

Conclusion

Paragon Protect & Restore supports various locations for backup data that should satisfy the requirements of the most demanding corporate data protection policies. The use of archiving to secondary storage instead of the pure parallel backup enables to avoid extra snapshots of ESX guests or physical Windows machines, to minimize backup windows, to release CPU and network bandwidth resources. PPR includes a very efficient backup data retention mechanism that can help guarantee storages are not overfilled and ready to take the most up-to-date backup or replica. Included Storage Browser allows easy and transparent management of backup data storages contain.

For additional information, please contact us at: **www.paragon-software.com**, **www.protect-restore.com** or **+1 (888) 347-5462**.