

Acronis[®] Backup & Recovery[™] 10 Management Server reports

Technical white paper

Table of Contents

1	Report data	3
2	Time format	4
3	Relationship between views	4
4	Relationship diagram	6
5	Current state reports	6
5.1	Machines	6
5.1.1	Description	6
5.1.2	Query samples	6
5.1.3	Column values	8
5.2	Tasks	9
5.2.1	Description	9
5.2.2	Query samples	9
5.2.3	Column values	10
5.2.4	List of task types	11
5.3	Policies	12
5.3.1	Description	12
5.3.2	Query samples	12
5.3.3	Column values	13
5.4	Backup plans	13
5.4.1	Description	13
5.4.2	Query samples	13
5.4.3	Column values	14
5.5	Archives and backups	15
5.5.1	Description	15
5.5.2	Query samples	15
5.5.3	Column values	16
5.5.4	Original data size	17
5.6	Vaults report	18
5.6.1	Description	18
5.6.2	Query samples	18
5.6.3	Column values	18
6	Historical reports	19
6.1	Vaults statistics report	19
6.1.1	Description	19
6.1.2	Query samples	19
6.1.3	Column values	20
6.2	Task activities report	21
6.2.1	Description	21
6.2.2	Query samples	21
6.2.3	Column values	21
Index		23

This document covers report views—a newly-introduced feature of the advanced editions of Acronis® Backup & Recovery™ 10.

These views are stored in a Microsoft SQL Server database which is used by Acronis Backup & Recovery 10 Management Server. Each view is similar to a database table.

By executing Structured Query Language (SQL) queries against these views, you can get a well-structured centralized overview of product performance and statistics throughout the enterprise.

Report views complement the reporting graphical user interface (GUI), which is now available in Acronis Backup & Recovery 10 Management Console when it is connected to the management server. An advantage of using the views is that you are free to run any custom queries—not necessarily covered by the reporting GUI—or to export the result to any format supported by Data Transformation Services (DTS).

The primary objective of this document is the description of the views and presentation of sample queries for getting some useful reports.

1 Report data

Where the report data is taken from

Reports are built based on the data that the management server retrieves from the machines that are registered on it. The data is retrieved periodically (normally, every 2 minutes), so there might be some minor latency between the actual data and the data in the reports.

The retrieved data is stored in the management server's databases.

Types of reports

There are two types of reports:

- Current
- Historical

Current reports reflect the current state of tasks, backup plans, backup policies, machines, archives, backups, and vaults. You can also view all this information in the Acronis Backup & Recovery 10 Management Console.

Historical reports are built based on the logs that the management server has collected from the registered machines over time. They provide information about vaults' statistics (for example, the history of the managed vaults' space usage) and task activities (for example, the result with which a task finished each time).

Historical reports enable you to obtain information that is not shown in the management console, such as how the free space in a managed vault changed within the reporting period.

SQL server instances used by the management server

The management server uses two instances of Microsoft SQL servers, called the operational SQL server instance and the reporting SQL server instance. Both instances are specified during the management server installation.

By default, the management server is installed together with Microsoft SQL Server 2005 Express, and the **ACRONIS** instance of this SQL server is used as both the operational and the reporting SQL server instance.

Accessing the report views

The report views (and necessary functions) for all reports are stored in the **acronis_reports** database on the reporting SQL server instance. You can access the report views and execute queries against them by using any management tool that works with that server instance.

For Microsoft SQL Server 2005 Express, you can use Microsoft SQL Server Management Studio Express—a free management tool available from the Microsoft Web site at:

<http://www.microsoft.com/downloads/details.aspx?FamilyID=c243a5ae-4bd1-4e3d-94b8-5a0f62bf7796&displaylang=en>

If the reporting SQL server instance and the management tool you are using are installed on different computers, you may need to configure the SQL server instance to allow remote connections. For Microsoft SQL Server 2005 Express, you will find the configuration instructions in the following article from the Microsoft Support Web site: <http://support.microsoft.com/kb/914277>

2 Time format

All dates and times in the report views are stored in the Coordinated Universal Time (UTC) format.

3 Relationship between views

Each entry in the report views has its unique identifier. Each view also contains columns that establish a relationship between dependent entries. For example, the **Tasks** view contains a column that identifies the backup plan to which a particular task belongs.

Machines

Information about all physical machines and virtual machines that can be centrally managed by Acronis Backup & Recovery 10 is stored in the **Machines** view.

A machine where at least one Acronis Backup & Recovery 10 agent is installed is considered as a physical machine.

A virtual machine is a machine that is managed by Agent for ESX/ESXi or Agent for Hyper-V.

A machine where Agent for ESX/ESXi is installed—which technically is a ESX/ESXi virtual appliance—is considered as a physical machine, according to the above definition of a physical machine.

The unique identifier of a machine is given in the **ManagedEntityID** column of the **Machines** view. (“Managed entity” is a term used internally in the product.)

The **HostID** column contains the unique identifier of the physical machine or of the virtual machine's host.

The **ManagedEntityID** column is also present in the **Tasks**, **BackupPlans**, and **ArchivesAndBackups** views. This enables you, for example, to determine the machine on which a given backup plan exists.

Backup plans, backup policies, and tasks

A backup plan consists of one or more tasks.

Backup plans can be local or centralized. A local backup plan is created on a machine directly, whereas a centralized backup plan is created as a result of deploying a backup policy.

Some tasks, such as recovery tasks and tasks running on Acronis Backup & Recovery 10 Storage Node, do not belong to any backup plan.

Information about all tasks, backup plans, and backup policies is contained in the **Tasks**, **BackupPlans**, and **Policies** views, respectively. The **ID** column in each view uniquely identifies the corresponding entry.

In the **Tasks** view, the **BackupPlanID** column stores the unique identifier of the backup plan to which the task belongs.

In the **BackupPlans** view, for centralized backup plans, the **PolicyID** column stores the identifier of the backup policy on which the backup plan is based.

Vaults, storage nodes, and archives and backups

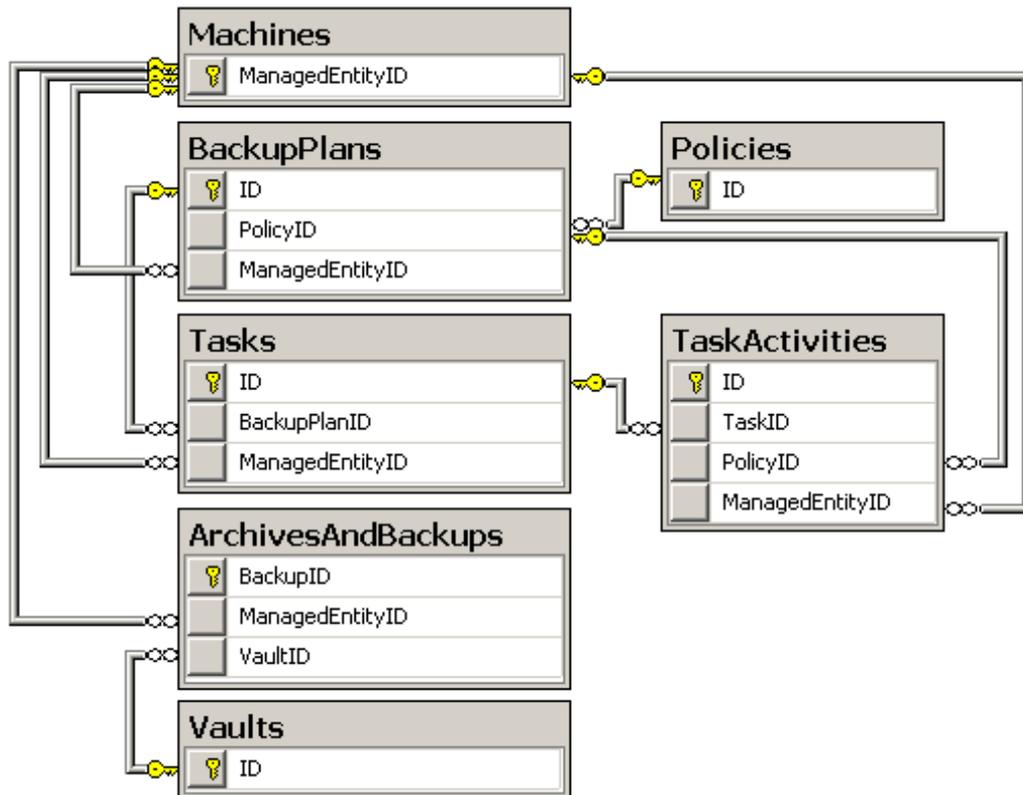
The **Vaults** view contains information about all centralized vaults—both managed and unmanaged. The **ID** column uniquely identifies the vault.

For a managed vault, the **HostID** column of this view stores the unique identifier of the storage node that manages the vault.

The **ArchivesAndBackups** view contains information about all archives and backups that are stored in the managed vaults. The **VaultID** column in this view identifies the managed vault in which the archive is stored.

4 Relationship diagram

The following diagram shows the relationship between the report views. Only the columns that uniquely identify the entry and establish the relationships are shown here.



5 Current state reports

5.1 Machines

5.1.1 Description

The **Machines** view represents all physical and virtual machines. You can differentiate between the types of machines by using the **ManagedEntityType** column.

5.1.2 Query samples

- The following query returns all “critical” physical and virtual machines—the ones with a last backup time or last connection time of more than five days, and the ones that have never been backed up:

```
SELECT * FROM Machines WHERE
(DATEDIFF(d, LastConnection, GETUTCDATE()) > 5 OR
DATEDIFF(d, LastBackup, GETUTCDATE()) > 5 OR
LastBackup IS NULL) AND
((ManagedEntityType = 'PHYSICAL MACHINE') OR
(ManagedEntityType = 'VIRTUAL MACHINE'))
```

Note that the function **GETUTCDATE** is used here instead of **GETDATE**, because all times in the databases are in the UTC format.

- The following query returns the name and the time of the last successful backup for all physical and virtual machines; the times are shown in the local time zone (that of the SQL server):

```
SELECT
    Name,
    ManagedEntityType,
    DATEADD(hh, DATEDIFF(hh, GETUTCDATE(), GETDATE()), LastBackup)
    AS [Last Successful Backup]
FROM Machines
WHERE
    ManagedEntityType = 'PHYSICAL MACHINE' OR
    ManagedEntityType = 'VIRTUAL MACHINE'
```

The following illustration shows the result as seen in Microsoft SQL Server Management Studio Express.

The screenshot shows the Microsoft SQL Server Management Studio Express interface. The query window displays the following SQL query:

```
SELECT
    Name,
    ManagedEntityType,
    DATEADD(hh, DATEDIFF(hh, GETUTCDATE(), GETDATE()), LastBackup)
    AS [Last Successful Backup]
FROM Machines
WHERE
    ManagedEntityType = 'PHYSICAL MACHINE' OR
    ManagedEntityType = 'VIRTUAL MACHINE'
```

Below the query window, the Results tab is active, showing a table with the following data:

	Name	ManagedEntityType	Last Successful Backup
1	10.250.148.64-VA	PHYSICAL MACHINE	NULL
2	BioHazard_Win2K3	VIRTUAL MACHINE	2009-10-24 00:49:27.000
3	WinXP	VIRTUAL MACHINE	2009-11-22 17:47:14.000
4	BioHazard_Win2K	VIRTUAL MACHINE	2009-11-30 23:39:39.000
5	TTTTT	VIRTUAL MACHINE	2009-11-30 21:42:23.000
6	AMS+LS_11345_EL	VIRTUAL MACHINE	2009-11-30 20:13:54.000
7	Ltest-new-w2k3	PHYSICAL MACHINE	2009-11-30 21:39:47.000
8	BioHazard_new2k	VIRTUAL MACHINE	2009-10-26 20:41:41.000
9	Windows 2003	VIRTUAL MACHINE	2009-10-29 07:03:32.000
10	10.250.144.212-VA	PHYSICAL MACHINE	NULL

The status bar at the bottom of the window shows: Item(s) Saved, Ln 5, Col 28, Ch 28, INS.

- The following query returns the list of all offline machines that run any of the Windows Server operating systems:

```
SELECT * FROM Machines WHERE
OperatingSystem LIKE '%Windows%Server%' AND
Availability = 'OFFLINE'
```

- The following query returns the list of virtual machines that are hosted on the ESX/ESXi host with IP address 10.1.2.3 and are managed by Agent for ESX/ESXi:

```
SELECT
  Machine.Name,
  Machine.Availability,
  Machine.LastBackup,
  Machine.Status
FROM Machines AS Machine
INNER JOIN Machines AS Appliance ON Machine.HostID = Appliance.HostID
WHERE
  Machine.ManagedEntityType = 'VIRTUAL MACHINE' AND
  Appliance.Name = '10.1.2.3-VA'
ORDER BY Machine.Name
```

5.1.3 Column values

The **Machines** view contains the following columns:

- **Name** – The name of the machine. This name is determined by the operating system that the machine runs.
- **Status** – The status of the machine, which depends on the statuses of the backup tasks that exist on it:
 - **OK** – All backup tasks are successful.
 - **WARNING** – There is at least one backup task that issued a warning during the backup, but there are no errors.
 - **ERROR** – At least one backup task has failed.
 - **UNKNOWN** – No backup tasks have yet run on the machine.
- **MachineAddress** – For a physical machine: the IP address of the physical machine. For a virtual machine: the IP address of the physical machine where the agent that manages the virtual machine is installed.
- **OperatingSystem** – The name of the operating system that the machine runs
- **LastConnection** – The date and time of the last synchronization between the machine and the management server
- **LastBackup** – The date and time when the latest *successful* backup finished on the machine
- **Availability** – The availability of the machine:
 - **ONLINE** – The latest connection to the machine by the management server was successful.
 - **OFFLINE** – The latest connection to the machine by the management server failed.
 - **WITHDRAWN** – The machine has been added to another management server or its management type has been set to stand-alone management.
 - **EXPIRED** – The management server failed to connect to the machine because the machine refused connection due to the expiration of a trial license.
- **ManagedEntityType** – The type of machine:
 - **PHYSICAL MACHINE** – The machine is a physical machine.
 - **VIRTUAL MACHINE** – The machine is a virtual machine.

- **VIRTUAL SERVER** – This value is reserved for use in future versions. Please disregard any rows that currently have this value.
 - **ManagedEntityID** – The unique identifier of the machine
 - **HostID** – For a physical machine: the unique identifier of the agent that is installed on the physical machine. For a virtual machine: the unique identifier of the agent that manages the virtual machine.
 - **HostType** – Reserved for use in future versions. Currently, the value in this column is always **AGENT**
 - **NextBackupStartTrigger** – The types of backup plan schedules on the machine:
 - **NONE** – There are no backup plans.
 - **MANUAL** – There are backup plans, but none of them run on a schedule (you can only start them manually).
 - **ON TIME** – There is a backup plan that runs on a time-based schedule—for example, daily.
 - **ON EVENT** – There is a backup plan that runs on an event-based schedule—for example, upon user logoff—and none of the backup plans run on a time-based schedule.
-
- Note: The backup plans that back up a virtual machine exist on the physical machine whose agent manages the virtual machine.*
-
- **NextBackupStartTime** – When the value in the **NextBackupStartTrigger** column is **ON TIME**, this column contains the time when the next backup that runs on a time-based schedule will start.

5.2 Tasks

5.2.1 Description

The **Tasks** view contains information about the tasks that exist on physical machines and on storage nodes. By using the **HostType** column, you can determine whether a task exists on a physical machine or on a storage node.

To retrieve more information about the physical machine on which the task exists, or about the backup plan to which the task belongs, you can join this view to the corresponding views by using the **LEFT JOIN** operator.

5.2.2 Query samples

- The following query returns the list of all failed backup tasks:

```
SELECT * FROM Tasks WHERE
LastResult = 'FAILED' AND
Type LIKE '%BACKUP%'
```

- The following query returns the list of registered machines on which no tasks have run within the past two days:

```
SELECT Machine FROM Tasks WHERE
HostType='AGENT'
GROUP BY HostID, Machine
HAVING MIN(DATEDIFF(d, LastFinishTime, GETUTCDATE())) > 2
```

5.2.3 Column values

The **Tasks** view contains the following columns:

- **Name** – The name of the task
- **Machine** – The name of the machine on which the task exists
- **ExecutionState** – The current execution state of the task:
 - **IDLE** – The task is idle.
 - **WAITING** – The task is ready to start, but is waiting until a condition is met or until another task running on the same machine has finished.
 - **RUNNING** – The task is in the process of execution.
 - **NEED INTERACTION** – The task needs user interaction to be able to proceed.
 - **STOPPING** – The task is in the process of stoppage.
- **Origin** – The origin of the task, which is determined by the type of the backup plan to which the task belongs:
 - **LOCAL STANDALONE** – The task does not belong to any backup plan.
 - **LOCAL** – The task belongs to a local backup plan.
 - **CENTRALIZED** – The task belongs to a centralized backup plan.
- **BackupPlan** – The name of the backup plan to which the task belongs—or **NULL** if the task does not belong to any backup plan
- **Type** – See List of task types (p. 11)
- **Schedule** – The type of the task's schedule:
 - **SCHEDULED** – The task is scheduled.
 - **MANUAL** – The task is not scheduled (you can only start it manually).
- **LastFinishTime** – The time when the task last finished, regardless of the last result
- **OwnerUID** – The unique identifier of the user who owns the task
- **Owner** – The name of the user who owns the task
- **ManagedEntityType** – For tasks that belong to a backup plan: the type of the machine whose identifier is given in the **ManagedEntityID** column:
 - **PHYSICAL MACHINE** – The machine is a physical machine.
 - **VIRTUAL MACHINE** – The machine is a virtual machine.For other tasks, the value in this column is: **UNKNOWN**
- **ManagedEntityID** – For tasks that belong to a backup plan: the unique identifier of the physical or virtual machine which the task backs up or whose backups it validates. For other tasks: **NULL**

Known issue: For tasks of a local backup plan that exists on an Agent for ESX/ESXi virtual appliance, the value in this column is the identifier of the virtual appliance.

- **Duration** – The amount of time, in seconds, that the task last ran, from start to finish.
- **LastResult** – The result of the latest task execution:
 - **NOT RUN YET** – The task has never run.
 - **SUCCEEDED** – The task completed successfully without warnings.
 - **SUCCEEDED WITH WARNINGS** – The task completed successfully with warnings.
 - **FAILED** – The task has failed.
 - **STOPPED** – The task has been stopped by the user or by the operating system (for example, due to the stoppage of the agent's service).

- **HostType** – The component that runs the task:
 - **AGENT** – Any Acronis Backup & Recovery 10 Agent
 - **STORAGE NODE** – Acronis Backup & Recovery 10 Storage Node
- **ID** – The unique identifier of the task
- **BackupPlanID** – The unique identifier of the backup plan to which the task belongs—or **NULL** if the task does not belong to any backup plan
- **HostID** – The unique identifier of the component that runs the task

***Note for Virtual Edition:** On a Hyper-V host, the value in this column will be the same for both the tasks performed by Agent for Hyper-V and by Agent for Windows. This is because Agent for Hyper-V is not a separate component, but a feature of Agent for Windows.*

5.2.4 List of task types

The following table lists the task types as they are named in the **Tasks** view and in the product’s graphical user interface.

<i>Name in the view</i>	<i>Name in the GUI</i>	<i>Task description</i>
ARCHIVE TO VM CONVERSION	Convert to VM	Converts a backup to a virtual machine
ARCHIVE VALIDATION	Validation (archive)	Validates an archive
ASRM ACTIVATION TASK	ASRM activation/deactivation	Activates or deactivates Acronis Startup Recovery Manager (ASRM)
ATTACHLOCATIONTASK	Vault attach	Attaches a managed vault to a different storage node
CLEAN UP	Cleanup	Cleans up an archive according to retention rules
CREATEASZTASK	ASZ creation	Creates Acronis Secure Zone (ASZ)
DELETE ARCHIVES TASK	Backup archives deletion	Deletes one or more archives
DELETE SLICES TASK	Backup deletion	Deletes one or more backups (internally called slices) from an archive
DELETEASZTASK	ASZ deletion	Deletes Acronis Secure Zone
DISK BACKUP	Backup (disk)	Performs a disk-level backup
DISK MANAGEMENT TASK	Disk management	Manages a disk or volume, for example changes a volume label
DISKMBRRESTORETASK	Recovery (MBR)	Recovers the master boot record (MBR) of a disk from a disk backup
DISKRESTORETASK	Recovery (disk)	Recovers disks from a disk backup; see also VOLUMERESTORETASK
EXPORT ARCHIVE	Export (archive)	Exports an archive
EXPORT SLICES	Export (backup)	Exports one or more backups of an archive
FILE BACKUP	Backup (file)	Performs a file-level backup

<i>Name in the view</i>	<i>Name in the GUI</i>	<i>Task description</i>
FILERESTORETASK	Recovery (file)	Recovers files from a file backup or from a disk backup
INVENTORYTASK	Inventorying	The inventorying task on a storage node: performs inventorying of tapes in a tape library
LOCATION VALIDATION	Validation (vault)	Validates a vault
MANAGEASZTASK	ASZ management	Manages Acronis Secure Zone—for example, changes its size
RESCANTASK	Rescan	The rescan task on a storage node: updates the database that contains information about tapes
RESTORE DISKS (EXISTING VM)	Recovery (disk to existing VM)	Recovers disks to an existing virtual machine (VM)
RESTORE DISKS (NEW VM)	Recovery (disk to new VM)	Recovers disks to a new virtual machine
RESTORE VOLUMES (EXISTING VM)	Recovery (existing VM)	Recovers volumes to an existing virtual machine
RESTORE VOLUMES (NEW VM)	Recovery (new VM)	Recovers volumes to a new virtual machine
SLICE VALIDATION	Validation (backup)	Validates a backup
STORAGE SERVER COMPACTING	Compacting	The compacting task on a storage node: removes unused data from deduplicated vaults
STORAGE SERVER LOCATION INDEXING	Indexing	The indexing task on a storage node: performs deduplication at target
VIRTUAL MACHINE BACKUP	Backup (virtual machine)	Backs up a virtual machine
VOLUMERESTORETASK	Recovery (volume)	Recovers volumes from a disk backup; see also DISKRESTORETASK

5.3 Policies

5.3.1 Description

The **Policies** view contains information about all backup policies. The type of backup policy is determined by the **SourceType** column.

5.3.2 Query samples

- The following query returns the backup policies that are not applied to any machine (and thus unused):

```
SELECT * FROM Policies WHERE
IsDeployed = 0 AND
IsDeploying = 0 AND
IsUpdating = 0 AND
IsRevoking = 0
```

- The following query returns failed backup policies that perform file-level backups:

```
SELECT * FROM Policies WHERE
Status = 'ERROR' AND
SourceType = 'CENTRALIZED FILE BACKUP'
```

5.3.3 Column values

The **Policies** view contains the following columns:

- **Name** – The name of the policy
- **SourceType** – The type of data that is backed up under the policy:
 - **CENTRALIZED DISK BACKUP** – Disks or volumes of physical or virtual machines
 - **CENTRALIZED FILE BACKUP** – Files
- **IsDeploying** – Equals 1 if the policy is in the **Deploying** state on at least one machine; otherwise, equals 0
- **IsDeployed** – Equals 1 if the policy is in the **Deployed** state on at least one machine; otherwise, equals 0
- **IsUpdating** – Equals 1 if the policy is in the **Updating** state on at least one machine; otherwise, equals 0
- **IsRevoking** – Equals 1 if the policy is in the **Revoking** state on at least one machine; otherwise, equals 0
- **Status** – The cumulative status of the policy statuses on all machines the policy is applied to:
 - **OK** – The policy status on all machines is **OK**.
 - **WARNING** – There is a machine on which the policy status is **Warning**, and there are no machines on which the policy status is **Error**.
 - **ERROR** – There is a machine on which the policy status is **Error**.
- **Schedule** – The schedule type of the policy:
 - **SCHEDULED** – The policy is scheduled.
 - **MANUAL** – The policy is not scheduled (you can only start the corresponding centralized backup plans manually).
- **Owner** – The name of the user who owns the policy
- **Comments** – Policy comments
- **OwnerUID** – The unique identifier of the user who owns the policy
- **ID** – The unique identifier of the policy

5.4 Backup plans

5.4.1 Description

The **BackupPlans** view contains the list of all local and centralized backup plans. The origin of the backup plan (local or centralized) is determined by the **Origin** column.

5.4.2 Query samples

- The following query returns the scheduled backup plans that have not run within the past week:

```
SELECT * FROM BackupPlans WHERE
ScheduleType = 'SCHEDULED' AND
DATEDIFF(d, LastBackupTime, GETUTCDATE()) > 7
```

- The following query returns all failed backup plans that back up virtual machines:

```
SELECT * FROM BackupPlans WHERE
ManagedEntityType = 'VIRTUAL MACHINE' AND
Status = 'ERROR'
```

5.4.3 Column values

The **BackupPlans** view contains the following columns:

- **Name** – The name of the backup plan. For centralized backup plans, this is the same as the name of the corresponding backup policy
- **Owner** – The name of the user who owns the backup plan
- **Type** – The type of data that is backed up under the backup plan:
 - **CENTRALIZED DISK BACKUP** – Disks or volumes of physical or virtual machines; centralized backup plan
 - **CENTRALIZED FILE BACKUP** – Files; centralized backup plan
 - **VIRTUAL DISK BACKUP** – Disks or volumes of a virtual machine; local backup plan
 - **VIRTUAL DISK GROUP BACKUP** – One or more entire virtual machines; local backup plan
 - **DISK BACKUP** – Disks or volumes of a physical machine; local backup plan
 - **FILE BACKUP** – Files of a physical machine; local backup plan

*Known issue: In build 11345, the value **VIRTUAL DISK BACKUP** is shown when the value **VIRTUAL DISK GROUP BACKUP** should be shown; and the value **UNKNOWN** is shown when the value **VIRTUAL DISK BACKUP** should be shown. This issue will be fixed in the next product updates.*

- **Origin** – The origin of the backup plan:
 - **LOCAL** – Local backup plan
 - **CENTRALIZED** – Centralized backup plan
- **ExecutionState** – The current execution state of the backup plan, which is based on the states of its tasks:
 - **NEED INTERACTION** – At least one task needs user interaction.
 - **RUNNING** – At least one task is running, and there are no tasks that need user interaction.
 - **WAITING** – At least one task is waiting, and there are no tasks that are running or need user interaction.
 - **STOPPING** – At least one task is stopping, and there are no tasks that are running, need user interaction, or waiting.
 - **IDLE** – All the tasks are idle.
- **Status** – The status of the backup plan, which is based on the statuses of its backup tasks:
 - **OK** – All the backup tasks have completed successfully.
 - **WARNING** – At least one backup task has succeeded with warnings, and there are no backup tasks that have failed.
 - **ERROR** – At least one backup task has failed.
- **ScheduleType** – The type of the backup plan's schedule:
 - **SCHEDULED** – The backup plan is scheduled.
 - **MANUAL** – The backup plan is not scheduled (you can only start its tasks manually).
- **LastBackupTime** – The date and time when the last backup under the backup plan finished, regardless of the result.

- **Comments** – Backup plan comments. For a centralized backup plan, the value in this column is the same as in the **Comments** column for the corresponding backup policy.
- **ManagedEntityType** – The type of the machine whose identifier is given in the **ManagedEntityID** column:
 - **PHYSICAL MACHINE** – The machine is a physical machine.
 - **VIRTUAL MACHINE** – The machine is a virtual machine.
- **OwnerUID** – The unique identifier of the user who owns the backup plan
- **ID** – The unique identifier of the backup plan
- **PolicyID** – For a centralized backup plan: the unique identifier of the backup policy on which the backup plan is based.

***Note:** For a local backup plan, this column identifies an equivalent of the backup policy, which uniquely corresponds to the backup plan. Through this column, the **BackupPlans** view is linked with the **TaskActivities** view (p. 21)—see the **PolicyID** column in that view.*

- **ManagedEntityID** – The unique identifier of the physical machine on which the backup plan exists
- **HostID** – The unique identifier of the agent that runs the backup plan's tasks

5.5 Archives and backups

5.5.1 Description

The **ArchivesAndBackups** view contains the list of all backups in all archives from all the *managed* vaults. It is possible to track the relationship of an archive to the managed vault.

Each entry in the view corresponds to a backup. Archive-specific columns, such as **ArchiveType** or **NumberOfFullBackups**, contain information about the archive in which the backup is contained; their content is the same for all backups in the archive.

Backups of virtual and physical machines are combined in this view. To differentiate them, you should join this view with the **Machines** view and filter the result by the machine type.

5.5.2 Query samples

- The following query returns the list of all full backups that are bigger than 10 GB:

```
SELECT * FROM ArchivesAndBackups WHERE
BackupDataSize > 10 * 1024 * 1024 * CAST(1024 AS BIGINT) AND
BackupType = 'FULL'
```

***Note:** The **CAST** function is used here because the value of 10 GB falls outside the **INT** type domain.*

- The following query returns the list of all archives that were created more than a year ago but do not contain any backups created within the past month:

```
SELECT * FROM ArchivesAndBackups A WHERE
DATEDIFF(yy, ArchiveCreationTime, GETUTCDATE()) > 1 AND
NOT EXISTS(SELECT * FROM ArchivesAndBackups WHERE DATEDIFF(mm,
BackupCreationTime, GETUTCDATE()) <= 1 AND ArchiveID = A.ArchiveID)
```

5.5.3 Column values

The **ArchivesAndBackups** view contains the following columns. In this description, “archive” refers to the archive that contains the given backup; “corresponding backup plan” refers to the backup plan under which the backup was created:

- **Name** – The name of the archive
- **ArchiveType** – The type of archive:
 - **FILE** – The archive contains file backups.
 - **IMAGE** – The archive contains disk backups.
- **ArchiveCreationTime** – The date and time when the archive was created. It does not change upon adding or deleting backups in the archive.
- **ArchiveOwnerName** – The name of the user who owns the archive. This is the user whose credentials are used to access the managed vault where the archive is stored.
- **MachineName** – The name of the machine on which the corresponding backup plan ran.
The following four columns show the different notions of size (in bytes) for the archive, which are calculated based on the corresponding sizes of the backups stored in it:
 - **ArchiveOccupiedSpace** – The sum of the values in the **BackupOccupiedSpace** column for all backups in the archive
 - **ArchiveBackedUpDataSize** – The sum of the values in the **BackupDataSize** column for all backups in the archive
 - **ArchiveUniqueDataSize** – The sum of the values in the **BackupUniqueDataSize** column for all backups in the archive
 - **ArchiveOriginalDataSize** – The sum of the values in the **BackupOriginalDataSize** column for all backups in the archive
- **NumberOfFullBackups** – The number of full backups that are stored in the archive
- **NumberOfIncrementalBackups** – The number of incremental backups that are stored in the archive
- **NumberOfDifferentialBackups** – The number of differential backups that are stored in the archive
- **ArchiveComments** – The archive comments: in the GUI, the contents of the **Archive comments** field of the backup plan under which the archive is created
- **IsEncrypted** – Equals 1 if the archive is encrypted; otherwise, equals 0
- **EncryptionAlgorithm** – Encryption level for the archive. Archives can be encrypted by using the Advanced Encryption Standard (AES), with a key size of 128, 192, or 256 bits.
 - **NONE** – The archive is not encrypted.
 - **AES128** – The archive is encrypted by using a 128-bit key.
 - **AES192** – The archive is encrypted by using a 192-bit key.
 - **AES256** – The archive is encrypted by using a 256-bit key.
- **ArchiveOwnerUID** – The unique identifier of the user who owns the archive
- **BackupType** – The type of backup:
 - **FULL** – Full backup
 - **INCREMENTAL** – Incremental backup
 - **DIFFERENTIAL** – Differential backup
- **BackupCreationTime** – The date and time when the backup was created

- **BackupComments** – The backup comments: in the GUI, the contents of the **Archive comments** field of the backup plan under which the backup was created
- **IsSectorBySectorBackup** – Equals 1 if the backup is sector-by-sector (as opposed to a backup performed by taking a disk snapshot); otherwise, equals 0
- **IsLiveLinuxBackup** – Equals 1 if the backup was created by Agent for Linux; otherwise, equals 0
- **BackupOccupiedSpace** – The size actually occupied by the backup
- **BackupUniqueDataSize** – The amount of backed up data that is not deduplicated or indexed at the moment. After indexing of the backup is completed, this data includes disk blocks of a nonstandard size (for disk backups) and small files (for file backups).
If deduplication is not performed on the backup (for example, when the vault is non-deduplicated or when the archive is encrypted), the value in this column is the same as in the **BackupDataSize** column.
- **BackupDataSize** – The amount of data that was backed up to the backup.
- **BackupOriginalDataSize** – The amount of data that can be recovered from the backup. For a full backup, this amount is the same as in the **BackupDataSize** column. For an incremental or differential backup, this amount may include data from the backups on which the given backup depends. You will find an example in the section Original data size (p. 17).
- **ArchiveID** – The unique identifier of the archive
- **VaultID** – The unique identifier of the managed vault in which the archive is stored
- **HostID** – The unique identifier of the agent that ran the corresponding backup plan.
- **ManagedEntityID** – The unique identifier of the physical machine where the corresponding backup plan ran.
- **BackupID** – The unique identifier of the backup

5.5.4 Original data size

The original data size of a backup shows the amount of data that can be recovered from that backup. To illustrate this size, let us suppose the following:

- You have an archive consisting of two backups: the first backup is full and the second backup is incremental.
- The amount of data that was backed up to the first and second backups is 100 GB and 10 GB, respectively.
- The 10 GB in the second backup include 5 GB of data that changed since the first backup, and 5 GB of new data.

In this case, the amount of data that you can recover from the first (full) backup is 100 GB.

The amount of data that you can recover from the second (incremental) backup is 105 GB; namely, 10 GB of the new and changed data that is stored in this backup, plus 95 GB of the unchanged data that is stored in the full backup.

Thus, the original data sizes of the first and second backups are 100 GB and 105 GB, respectively.

The original data size of an archive is the sum of the original data sizes of its backups—in this example, 205 GB. It shows the total size of the sets of data that can be recovered from the archive.

The original data size of a vault is the sum of the original data sizes of the archives stored in the vault.

5.6 Vaults report

5.6.1 Description

The **Vaults** view contains the centralized vaults—both managed and unmanaged. The vault type (whether the vault is managed or not) is determined by the **IsManaged** column. The **HostID** column refers to the storage node that currently manages the vault (for an unmanaged vault, the value in this column is **NULL**).

5.6.2 Query samples

- The following query returns the managed vaults whose storage is located on a network share with a name starting with **\\server\backups**

```
SELECT * FROM Vaults WHERE  
IsManaged = 1 AND  
StoragePath LIKE '\\server\backups%'
```

- The following query returns the managed vaults that are both deduplicated and compressed—and thus are potentially slow:

```
SELECT * FROM Vaults WHERE  
IsDeduplicated = 1 AND  
CompressionLevel <> 'NONE'
```

5.6.3 Column values

The **Vaults** view contains the following columns:

- **Name** – The name of the vault
- **IsManaged** – Equals 1 if the vault is managed; otherwise, equals 0
- **OwnerName** – The name of the user who created the vault
- **OwnerUID** – The unique identifier of the user who created the vault
- **Comments** – The vault comments
- **StoragePath** – The path to the folder where the archives are stored and—for deduplicated vaults—where the deduplication data store is located
- **DatabasePath** – The path to the folder where the vault's database is located. For centralized unmanaged vaults, the value in this column is an empty string (' ').
- **IsDeduplicated** – Equals 1 if the vault is deduplicated; otherwise, equals 0
- **EncryptionAlgorithm** – Encryption level for the vault. Managed vaults can be encrypted by using the Advanced Encryption Standard (AES), with a key size of 128, 192, or 256 bits. Unmanaged vaults cannot be encrypted.
 - **NONE** – The vault is not encrypted.
 - **AES128** – The managed vault is encrypted by using a 128-bit key.
 - **AES192** – The managed vault is encrypted by using a 192-bit key.
 - **AES256** – The managed vault is encrypted by using a 256-bit key.
- **CompressionLevel** – The compression level for the vault. Only deduplicated vaults can be compressed.
 - **NONE** – The vault is not compressed.
 - **NORMAL** – The deduplicated vault is compressed.
- **ID** – The unique identifier of the vault

- **HostID** – The unique identifier of the storage node that manages the vault. For an unmanaged vault, the value in this column is **NULL**.

6 Historical reports

6.1 Vaults statistics report

6.1.1 Description

The **GetVaultsStatistics** function returns a table containing the statistics for all managed vaults that existed within a period you specify. The **VaultID** column in this table identifies the vault whose statistics are presented.

Because this is a historical report, the table may include vaults that no longer exist. By using this table together with the **Vaults** view, you can determine the missing vaults.

The **GetVaultsStatistics** function takes three arguments: the start of the reporting period; the end of the reporting period; and the granularity of the report (the interval between consecutive points in time), in seconds.

For example, the following query will return vaults' statistics as of 01:00, 02:00, and 03:00 on January 1, 2010 (UTC):

```
SELECT *
FROM dbo.GetVaultsStatistics('1-1-2010 01:00:00.000', '1-1-2010 03:00:00.000',
3600)
```

The management server updates statistics about the managed vaults every 15 minutes (900 seconds), so specifying a finer granularity will not increase report precision.

6.1.2 Query samples

- The following query returns vault statistics for the past two months on a daily basis:

```
SELECT * FROM dbo.GetVaultsStatistics(DATEADD(m, -2, GETUTCDATE()),
GETUTCDATE(), 86400)
```

- The following query returns the periods when any vault on the storage node named **datastore** was unavailable (or the management server was down) within the past week:

```
SELECT * FROM dbo.GetVaultsStatistics(DATEADD(d, -7, GETUTCDATE()),
GETUTCDATE(), 900) WHERE
NumberOfArchives IS NULL AND
StorageNodeName = 'datastore'
```

- The following query returns the average total ratio among all the deduplicated vaults of the **datastore** storage node on a daily basis for a month:

```
SELECT AVG(TotalRatio) FROM dbo.GetVaultsStatistics(DATEADD(m, -1,
GETUTCDATE()), GETUTCDATE(), 86400) WHERE
StorageNodeName = 'datastore' AND
IsDeduplicated = 1
GROUP BY Date
ORDER BY Date
```

6.1.3 Column values

The **GetVaultsStatistics** function returns a table that contains the following columns:

- **Date** – The point in time for which the statistics are presented
- **FreeSpace** – The amount of free space in the managed vault—normally, the amount of free space on the device that stores the managed vault
- **OccupiedSpace** – The space occupied by all archives that are stored in the managed vault, plus the size of the vault database and—for deduplicated vaults—the size of the deduplication data store
- **NumberOfArchives** – The number of archives that are stored in the managed vault
- **NumberOfBackups** – The total number of backups that are contained within the archives in the managed vault

The following four columns contain information about the size of the data in the archives, in bytes. It is calculated based on the archive sizes shown in the **ArchivesAndBackups** view (p. 15).

- **OriginalDataSize** – The sum of the values in the **ArchiveOriginalDataSize** column for all archives that are stored in the managed vault
- **ArchiveDataSize** – The sum of the values in the **ArchiveBackedUpDataSize** column for all archives that are stored in the managed vault
- **UniqueDataSize** – The sum of the values in the **ArchiveUniqueDataSize** column for all archives that are stored in the managed vault. For a non-deduplicated vault, the value in this column is **NULL**.
- **DeduplicatedArchiveSize** – The amount of data in the archives that was removed as a result of deduplication. For a non-deduplicated vault, the value in this column is 0. For a deduplicated vault, the value in this column is the value in **ArchiveDataSize** minus the value in **UniqueDataSize**.

The following four columns contain information about data ratios for the contents of the managed vault. Smaller ratios mean higher storage efficiency.

These ratios are calculated based on the sizes shown in the previous four columns.

- **TotalRatio** – The value in **OccupiedSpace** divided by the value in **OriginalDataSize**. This ratio shows the space actually occupied by the archives as compared with the overall size of the sets of data that can be recovered from these archives. For an example on original data size, see Original data size (p. 17).
- **BackedUpDataRatio** – The value in **ArchiveDataSize** divided by the value in **OriginalDataSize**. This ratio shows how the use of incremental and differential backups has reduced the space needed to store the sets of data in the archives.
- **DeduplicationRatio** – The value in **UniqueDataSize** divided by the value in **ArchiveDataSize**. This ratio shows how deduplication alone has reduced the space needed for the archives. For a non-deduplicated vault, the value in this column is **NULL**.
- **CompressionRatio** – The value in **OccupiedSpace** divided by the value in **ArchiveDataSize**. This ratio shows how deduplication (for deduplicated vaults) and archive compression have reduced the space needed for the archives.
- **VaultName** – The name of the managed vault
- **VaultID** – The unique identifier of the managed vault
- **StorageNodeName** – The name of the computer on which the storage node that manages the vault is installed
- **StorageURI** – The same as the **StoragePath** column of the **Vaults** view
- **IsEncrypted** – Equals 1 if the managed vault is encrypted; otherwise, equals 0

- **IsDeduplicated** – Equals 1 if the managed vault is deduplicated; otherwise, equals 0

6.2 Task activities report

6.2.1 Description

Each run of a task, from start to finish, corresponds to a single task activity. Task activities are represented by the **TaskActivities** view containing all the task activities occurred during the period that can be deduced from logs. The retention of logs reduces the amount of information available to this view. In case the logs are corrupted, this view may be inconsistent.

By joining this view and the view of a current report, you can get a more thorough report. For example, by using the **PolicyID** column, you can get information about both the task activities and their backup plans.

Keep in mind, however, that, because this is the view of a historical report, it may include tasks which, together with the corresponding backup plans, no longer exist. This means that such backup plans are not present in the report views for current reports.

6.2.2 Query samples

- The following query returns all the tasks that failed on the machine **webserver.example.com** within the past week:

```
SELECT * FROM TaskActivities WHERE
DATEDIFF(m, FinishTime, GETUTCDATE()) < 7 AND
Machine = 'webserver.example.com' AND
Result = 'FAILED' ORDER BY StartTime
```

- The following query returns the failure/success statistics of all tasks for the last 30 days:

```
SELECT
    TaskID,
    TaskName,
    SUM(CASE WHEN Result = 'SUCCEEDED' THEN 1 ELSE 0 END) SuccessCount,
    SUM(CASE WHEN Result = 'SUCCEEDED WITH WARNINGS' THEN 1 ELSE 0 END)
WarningCount,
    SUM(CASE WHEN Result = 'FAILED' THEN 1 ELSE 0 END) FailureCount
FROM
    TaskActivities
WHERE
    DATEDIFF(d, FinishTime, GETUTCDATE()) < 30
GROUP BY TaskID, TaskName
ORDER BY TaskName ASC
```

6.2.3 Column values

The **TaskActivities** view contains the following columns:

- **TaskName** – The same as the **Name** column of the **Tasks** view (p. 10)
- **Owner** – The same as the **Owner** column of the **Tasks** view
- **ManagedEntityName** – For tasks that belong to a backup plan: the name of the physical or virtual machine which the task backs up or whose backups it validates. For other types of tasks: an empty string (' ').

Known issue: For tasks of a local backup plan that exists on an Agent for ESX/ESXi virtual appliance, the value in this column is the name of the virtual appliance.

- **ManagedEntityType** – The same as the **ManagedEntityType** column of the **Tasks** view
 - **BackupPolicy** – The name of the backup policy (for centralized tasks) or the local backup plan (for local tasks) on which the task is based. If the task does not belong to a backup plan, the value in this column is an empty string (' ').
 - **Machine** – The name of the physical machine on which the task exists.
For tasks belonging to a backup plan that backs up virtual machines, the value in this column is the name of the physical machine where the agent managing the virtual machine (such as Agent for ESX/ESXi) is installed, whereas the value in the **ManagedEntityName** column is the name of the virtual machine.
 - **StartTime** – The date and time when the task activity started
 - **FinishTime** – The date and time when the task activity finished, regardless of the result
-
- Tip: Based on the values in the **FinishTime** and **StartTime** columns, you can determine the duration of a task activity, including that of the latest one. The duration of the latest task activity is also stored in the **Duration** column of the **Tasks** view. These two values might differ by a few seconds due to logging latency.*
-
- **Result** – The result of the task activity. The possible values here are the same as those in the **LastResult** column of the **Tasks** view.
 - **ID** – The unique identifier of the task activity
 - **TaskID** – The same as the **ID** column of the **Tasks** view.
 - **ManagedEntityID** – The same as the **ManagedEntityID** column of the **Tasks** view, except that the identifier **00000000-0000-0000-0000-000000000000** is used here instead of **NULL**
 - **PolicyID** – The unique identifier of the backup policy on which the task is based. If the task is not based on a backup plan, the value in this column is **00000000-0000-0000-0000-000000000000**.
-
- Note: For a local backup plan, this column identifies an equivalent of the backup policy, which uniquely corresponds to the backup plan. Through this column, the **TaskActivities** view is linked with the **BackupPlans** view (p. 14)—see the **PolicyID** column in that view.*
-
- **HostID** – The same as the **HostID** column of the **Tasks** view

Index

A

Archives and backups • 15, 20

B

Backup plans • 13

C

Column values • 8, 10, 13, 14, 15, 16, 18, 20, 21, 22

Current state reports • 6

D

Description • 6, 9, 12, 13, 15, 18, 19, 21

H

Historical reports • 19

L

List of task types • 10, 11

M

Machines • 6

O

Original data size • 17, 20

P

Policies • 12

Q

Query samples • 6, 9, 12, 13, 15, 18, 19, 21

R

Relationship between views • 4

Relationship diagram • 6

Report data • 3

T

Task activities report • 21

Tasks • 9

Time format • 4

V

Vaults report • 18

Vaults statistics report • 19