

# **NORTHERN STORAGE CHARGEBACK USER'S GUIDE**



**NORTHERN**



Northern Storage Chargeback for Windows 2000/XP/2003 User Guide revision 2005.0.612.11

Copyright © 2005-2007 by Northern Parklife AB. All rights reserved.

NORTHERN is a registered trademark of Northern Parklife AB. Quota Server, Northern Quota Server, Northern Storage Suite, Northern Storage Reporter, Northern Storage Portal, Northern Storage Chargeback, Northern Storage Assistant, Northern Utilities, Northern Site Utilities, Optimover, RasTracker, PrintControl, MobiServer are trademarks of Northern Parklife AB.

Microsoft is a registered trademark of the Microsoft Corporation. Windows, Windows NT, Windows 2000, Windows XP and Windows 2003 are trademarks of Microsoft Corporation. All other trademarks are property of their respective owner.

NetApp and the Network Appliance logo are registered trademarks. Network Appliance and Manage ONTAP are trademarks of Network Appliance Inc. in the US and other countries.



# TABLE OF CONTENT

## INTRODUCTION

<b>Welcome to Northern Storage Chargeback .....</b>	<b>7</b>
<b>Concepts .....</b>	<b>7</b>
Chargeback Jobs .....	7
Billing Classes .....	8
Projections .....	8
Cost Centers .....	8
Data sources .....	9
Hosts .....	9
Credits .....	9
<b>User Interface .....</b>	<b>9</b>
Views .....	9

## CONFIGURATION

<b>Configuring Storage Chargeback .....</b>	<b>11</b>
<b>Overview .....</b>	<b>11</b>
<b>Installation .....</b>	<b>11</b>
<b>System Requirements .....</b>	<b>11</b>
<b>Licensing .....</b>	<b>11</b>
<b>Configuring the database connection .....</b>	<b>12</b>
Creating the database .....	12
Setting database permissions .....	12
Connecting Storage Chargeback to the database .....	12
<b>Adding a Host .....</b>	<b>13</b>
<b>Using Storage Chargeback .....</b>	<b>16</b>
<b>Configuring for MS Exchange Jobs .....</b>	<b>16</b>
<b>Licensing .....</b>	<b>18</b>
.....	18

## USAGE

<b>Chargeback Jobs .....</b>	<b>19</b>
Adding a Chargeback Job .....	19
Removing a Chargeback Job .....	25
Editing a Chargeback Job .....	25
Viewing a Chargeback Job result .....	26
Tracking Chargeback Jobs .....	28

<b>Billing classes</b> .....	<b>28</b>
About multi-tiered pricing models .....	29
About The Default Billing Class .....	30
Adding a Billing Class .....	30
Removing a Billing Class .....	33
Editing a Billing Class .....	34
<b>Projections</b> .....	<b>34</b>
about projections .....	34
Adding a Projection .....	35
Removing a Projection .....	39
Editing a Projection Definition .....	39
Viewing a Projection Result .....	40
Defining a Data Range .....	41
<b>Using Projections</b> .....	<b>44</b>
Fastest Growing Cost Centers .....	44
Cost Prognosis per Cost Center .....	46
End-of-Budget Prognosis .....	47
Accumulate Projections .....	48
Cost Prognosis per Cost Center (Accumulated) .....	49
End-of-Budget Prognosis (Accumulated) .....	51
.....	51

## **TECHNICAL**

<b>SQL Database Structure</b> .....	<b>53</b>
Concepts .....	53
Table Definitions .....	53
<b>About Projection Extrapolation</b> .....	<b>66</b>
<b>About MSDE</b> .....	<b>69</b>
<b>Configuring MS SQL Server</b> .....	<b>69</b>
Creating a new database in MS SQL .....	69
Setting database permissions in MS SQL .....	70

## **NORTHERN**

<b>Northern: The Company</b> .....	<b>73</b>
Business Concept .....	73
The Market .....	73
Product Development .....	74
<b>Customer Support Program</b> .....	<b>74</b>
Features .....	74
<b>Contacting Northern</b> .....	<b>75</b>
Street Address .....	75
Phone .....	75
Fax .....	75
E-mail .....	75
Web .....	75

# INTRODUCTION

## WELCOME TO NORTHERN STORAGE CHARGEBACK

Northern Storage Chargeback makes business sense out of senseless storage use. By metering storage use and calculating the attendant operating costs, Storage Chargeback can bring the cost of managing storage back to the responsible party. Storage Chargeback can collect usage data from any disk object [drive, share, folder], use quota usage information from Northern Quota Server, or mailbox usage information from MS Exchange. Northern Storage Chargeback is a complete, end-to-end solution that collects usage statistics, determines appropriate pricing models, and generates reports.

## CONCEPTS

The structure of information in Northern Storage Chargeback revolves around the following key concepts:

- Chargeback Jobs
- Billing Classes
- Cost Center
- Data sources
- Projections
- Hosts
- Credits

## CHARGEBACK JOBS

A Chargeback Job is a series of operations that collect usage data, convert this data into a cost and store the results in a database.

### Definition

A Chargeback Job definition includes all the attributes regarding how, where and when the job should be performed. All the attributes are user-defined through the Chargeback Job Wizard and include for instance:

- Scheduling
- Notification
- Data source and data location
- Cost Center
- Binding between Billing Class and Data Location

### Results

A Chargeback Job has one definition but many results. Each time the job is run, according to its definition it will produce a set of cost data - this is the Chargeback Job result. If the Chargeback Job definition is executed multiple times, it will generate multiple results.

The Chargeback Job result is roughly a list of calculated costs by Cost Center.

## **BILLING CLASSES**

In Northern Storage Chargeback you can define a variety of different pricing or cost patterns, called Billing Classes. A Billing Class can be used to define multi-tiered pricing models. The Billing Classes can be given descriptive names and you can assign Billing Classes arbitrarily to different locations [volume, path, quota or equivalent].

## **PROJECTIONS**

A Projection is a report that provides cost trend and analysis according to historical cost data. The past data in this case are the results of previously run Chargeback Jobs.

Northern Storage Chargeback proposes several types of Projection depending on what information you want to extract (Cost Prognosis, End-of-Budget Analysis and Fastest Growing Cost Center).

## **COST CENTERS**

Cost Center is a concept used in Northern Storage Chargeback to signify where usage and the associated costs originate. A Cost Center may thus be a user or a user group, it may also be a specific location [drive, path, folder], or a quota.

Which kind of Cost Center you wish to use is dependant on the aims of your Chargeback Job. Either the costs are carried by an individual or they should be carried by a collective unit, such as a department or a user group.

Cost Centers:

- Users
- User Groups
- Paths
- Quotas
- MS Exchange mailboxes

In Storage Chargeback User always refers to the Windows concept of user account. Group always refers to the Windows concept of primary user group.

It is important to notice that when you run a Chargeback Job on a certain server [running it for instance on c:\ on ALPHA], you will get highly differing results depending on whether you use User Group or User as your Cost Center.

If you specify User Group as your Cost Center the total collective usage of all members of each group will be compared to the relevant Billing Class, and the final cost will be calculated.

If you specify Users as your Cost Center, each user in turn will be compared to the relevant Billing Class, and the final cost will be calculated per individual account.

Paths, quotas or MS Exchange mailboxes are always resolved individually: each quota, path or MS Exchange mailbox is compared to the relevant Billing Class and the final cost is calculated per path, quota or MS Exchange mailbox.

## DATA SOURCES

You can collect storage usage statistics from three kinds of data source in your Chargeback Jobs:

- File system
- Quota size
- MS Exchange

### File System

A file system Chargeback Job will collect the size of different file system objects under one or several paths on a server. Northern Storage Chargeback will also determine ownership of the objects according to the chosen Cost Center and calculate costs using the specified Billing Classes.

### Northern Quota Server

The usage data collected from Northern Quota Server relates to the maximum allocated amount of storage for different users, groups or folders, i.e. their quota size. When you define a Chargeback Job based on Quota Server quotas, Northern Storage Chargeback queries an administrator-defined Quota Server server about the size of all quotas on that particular server.

### MS Exchange

Mailbox usage can be collected by Northern Storage Chargeback to be used as a data source. A Chargeback Job towards MS Exchange will retrieve the current usage of all mailboxes on an MS Exchange server and process them through the Billing Class to return a list of cost per mailbox.

## HOSTS

In Northern Storage Chargeback the basic unit of management is a Host. A Host is a server or a NAS running a supported environment and a copy of the Northern Storage Suite installation.

For optimal performance, Storage Chargeback has to be installed on every Host where you wish to run Chargeback Jobs. However, you can perform certain types of jobs against a server which does not have Northern Storage Suite installed. Before you can define a Chargeback Job for a Host you have to add that Host to the Northern Storage Suite configuration. See "Adding a Host" on page 13.

## CREDITS

In Northern Storage Chargeback costs are expressed in credits. A credit is any monetary unit that you wish to use, whether it be US dollars, Euros or Yen, etc.

## USER INTERFACE

### VIEWS

Northern Storage Chargeback provides multiple views, all displayed within Internet Explorer. Each view offers access to the main concept of Storage Chargeback.

### Jobs

The Jobs view displays all available Chargeback Jobs. It provides all operations for managing Chargeback Jobs, i.e. creating, deleting, editing or scheduling.

**Results**

The Results view displays all the Chargeback Job results. Each entry in this view is a summary of the result of a Chargeback Job. You can double-click each Chargeback Job result summary to drill down to a presentation of the full Chargeback Job result.

**Projections**

The Projections view displays all available Projection types. From this view, you can go to the Schedule or History pages of each specific Projection type.

**Projection's Schedule**

For each Projection type, the Schedule page allows you to create, modify or delete your Projection definitions. All those operations are done through the Chargeback Projection Wizard.

**Projection's History**

For each Projection type, the History page presents all previous runs of the Projection and the associated data. This data is kept in an SQL database until destroyed on user request.

# CONFIGURATION

## CONFIGURING STORAGE CHARGEBACK

To be able to use the core features of Northern Storage Chargeback, you will need to install the software on the target Host and connect it to an SQL database. It is possible to run Chargeback jobs on a server where Northern Storage Suite is not installed, but performance will most likely be slower.

When this one time configuration is done, you can then define a first Chargeback Job. Later on, you can improve the cost calculation by defining more precise Billing Classes.

Finally, after several Chargeback Jobs have reported their results, you can define a projection from those Chargeback Jobs to analyse how your storage cost will evolve in the future.

## OVERVIEW

To configure Northern Storage Chargeback you must make sure to take all the following steps

- Install the software on the target Host
- Connect the installation to a SQL database, see “Configuring the database connection” on page 12
- Configure the Host, as described in “Adding a Host” on page 13

To perform MS Exchange jobs you will also need to:

- Connect Storage Chargeback to MS Exchange. See “Configuring for MS Exchange Jobs” on page 16

To perform Quota Server jobs you will also need to install Northern Quota Server, as described in the Northern Quota Server User Guide.

## INSTALLATION

For in depth installation and deployment information please refer to the Deployment Guide, found in the Northern Storage Suite installation directory.

## SYSTEM REQUIREMENTS

For a list of the system requirements please refer to the Deployment Guide, found in the Northern Storage Suite installation directory.

## LICENSING

For in depth licensing information please refer to the Deployment Guide, found in the Northern Storage Suite installation directory.

## **CONFIGURING THE DATABASE CONNECTION**

Northern Storage Chargeback stores all collected and calculated data into an SQL database. Most of the configuration operations are done automatically by Storage Chargeback. The only steps you need to worry about are:

- Create the database in the DBMS
- Set access level on the created database
- Add a database connection in the Database Connection Configurator

### **CREATING THE DATABASE**

If you chose to use MSDE when installing the software, this step has already been done automatically for you.

Regardless if you chose to use MSDE or MS SQL, you can let the Northern Storage Suite installer create the database for you. Alternatively you can refer to “Creating a new database in MS SQL” on page 69 to make one from the MS SQL manager.

### **SETTING DATABASE PERMISSIONS**

If you chose to use MSDE when installing the software, this step has already been done automatically for you.

If you chose to use MS SQL you need to make sure that the NSS service is read and write access to its database. This can be achieved in 2 ways:

- specifying the connection account when connecting Storage Chargeback to the database
- making the NSS Service account owner of the created MS SQL database, as described in “Setting database permissions in MS SQL” on page 70.

### **CONNECTING STORAGE CHARGEBACK TO THE DATABASE**

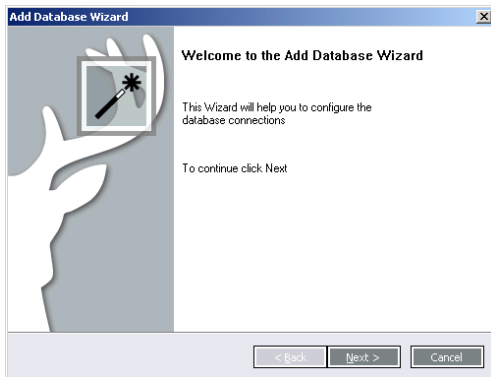
If, during the installation, you chose to install MSDE or to connect to an existing database (MSDE or MS SQL), this step has already been done automatically for you.

If you chose to use MS SQL and have created and configured the database in MS SQL as described above, you can now connect it to Northern Storage Chargeback. This is done in the Database Connection Configurator with the following steps:

- 1 On the Config page click the Database Connections link. The Database Connections control opens

2 Click the Add button. The Add Database Wizard appears.

Figure 1: Add Database Wizard



3 Follow the on-screen instructions.

The Master Server is a Storage Chargeback server which retains global informations (for example Projection definitions). Standard Servers do not require the Master Server to run normal operations.

## ADDING A HOST

Adding a Host merely enables Northern Storage Chargeback on that Host. It also writes the Host details into the database to be accessible from anywhere in the product. This is why the database connection has to be up and running before you can configure the host.

The configuration steps are very simple and are all done from the Storage Chargeback Web interface.

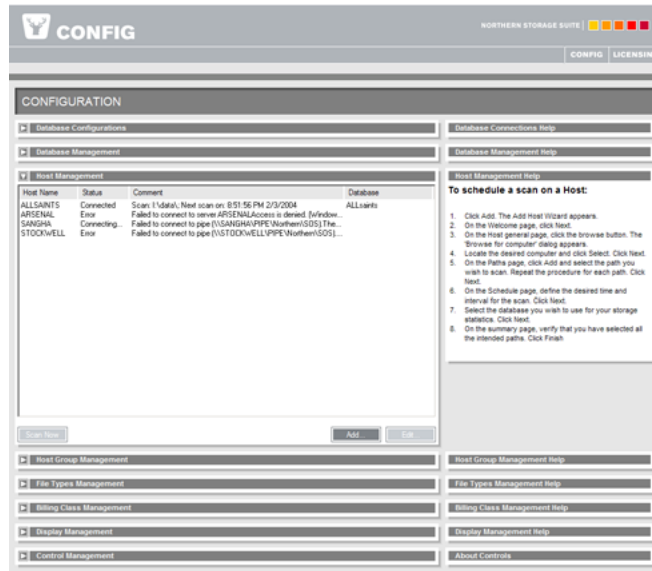


*Notice! If you want to add more database to be able to use several database servers, repeat the procedure as described. A second, third, fourth, etc. database connection will be added to the Database Connections List.*

To add a host:

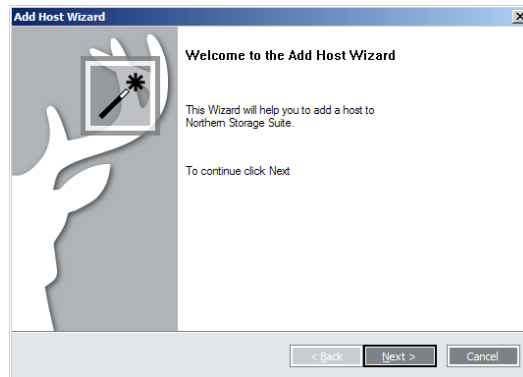
- 1 On the Config page, click the Hosts Management link. The Host Management Control opens

Figure 2: Host Management Configurator:



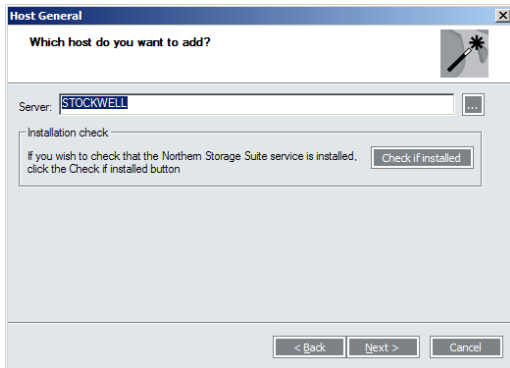
- 2 In the Hosts control, click the Add button. The Host Wizard appears.

Figure 3: Add Host Wizard



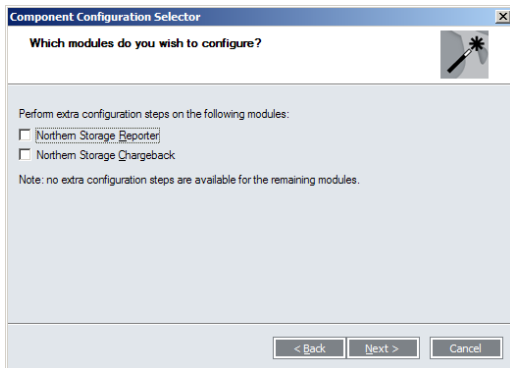
**3** In the 'Server' field, enter the name of the Host you wish to monitor. Click Next.

Figure 4: Server identification page



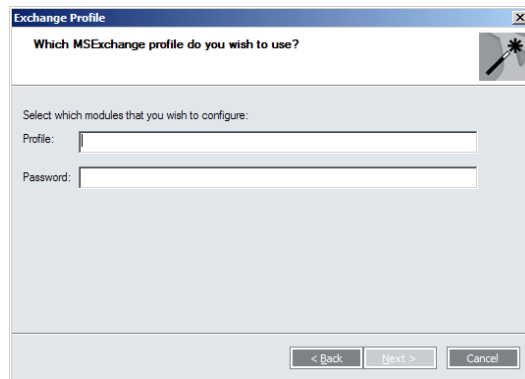
**4** On the Component Configuration Selector page, make sure to check "Storage Chargeback", then click Next.

Figure 5: Component Configuration Selector page



- 5 Enter the name and password for the MS Exchange profile. This step is required only if you plan to run Chargeback Job against MS Exchange. Note that the profile must be created separately as described in “Configuring for MS Exchange Jobs” on page 16. Leave the fields blank if you do not intend to use MS Exchange functionality. Click Next.

Figure 6: MS Exchange profile page



- 6 Check the summary on the Finish page and click Finish.

The Host is now added to the database and activated for Storage Chargeback operations. The only task left now is to define the Chargeback Jobs.

## USING STORAGE CHARGEBACK

After the one-time configuration you can now use Northern Storage Chargeback to analyse the cost of your storage. The Usage section details how to tell Storage Chargeback what you want it to do in the three main sections:

- See “Chargeback Jobs” on page 19, for creating new Chargeback Jobs
- See “Billing classes” on page 28, for defining new Billing Class
- See “Projections” on page 34, for adding new Projection

## CONFIGURING FOR MS EXCHANGE JOBS

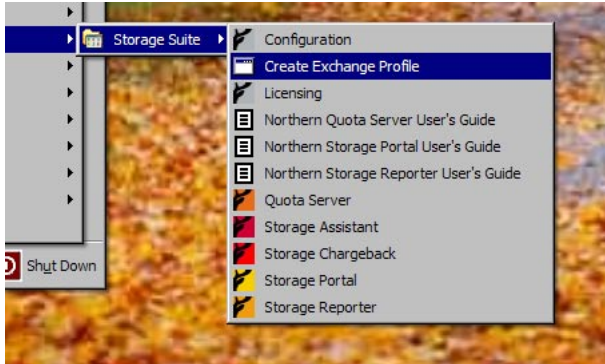
Due to the security features of MS Exchange, extra configuration steps are necessary to enable Chargeback Jobs toward MS Exchange.

The goal is to make MS Exchange allow Northern Storage Chargeback to query mailboxes information. This requires that the Northern Storage Suite is installed on the MS Exchange server and that the following procedure is performed when running the same account as the Northern Storage Suite service.

The service has to be installed on the MS Exchange server because MS Exchange profiles are local and cannot be accessed remotely. This is a security feature of MS Exchange to prevent remote illegal access to the sensitive contents of the mailbox.


To create the MS Exchange profile just run the "Create Exchange Profile" tool installed by the Northern Storage Suite setup.

Figure 7: Create Exchange Profile Start Menu items



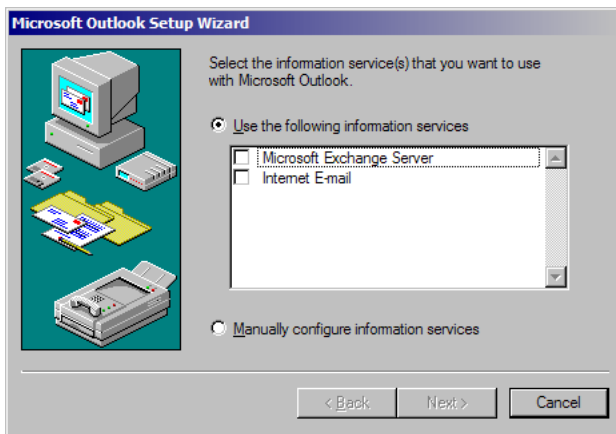
This will bring up a MAPI dialog with the list of all available profiles. You can create a new profile by clicking the "New..." button.


Figure 8: MAPI dialog for managing profiles



In the Microsoft Outlook Setup Wizard, you can then define a regular MS Exchange profile. Note that you do not need MS Outlook installed to complete the procedure.

Figure 9: Microsoft Outlook Setup Wizard



 *Notice! By default MS Exchange logs everytime a mailbox is accessed by a user different than the mailbox owner. With Storage Chargeback this will result in a lot of event log entries. Refer to MS Exchange documentation to turn off this logging..*

## **LICENSING**

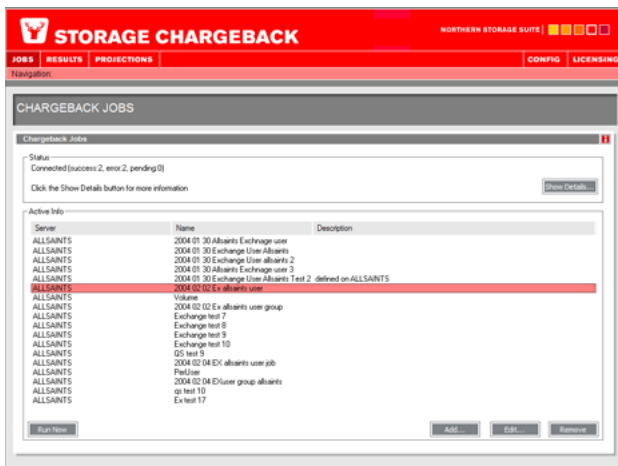
For in depth licensing information please refer to the Deployment Guide, found in the Northern Storage Suite installation directory.

# USAGE

## CHARGEBACK JOBS

Chargeback Job are managed from the Jobs View

Figure 10: The Storage Chargeback Jobs View



A Chargeback Job is based on three main elements. First you have to collect usage data from one of the three data sources [file system usage, Northern Quota Server's quota usage information, or MS Exchange's mailbox usage information]. Then you have to tell Storage Chargeback which Cost Center it should assign the usage data to [User, User Group, Location] When usage data has been collected and assigned it must be processed through user-defined pricing schemes.

The final results, i.e. the cost data, are stored into the database.

In more details a Chargeback Job has the following attributes:

- Data source
- Cost Center
- Pricing scheme [Billing Class]
- Scheduling
- Notification

Each attribute is a page of the Chargeback Job Wizard and is described in detail in the next section

### ADDING A CHARGEBACK JOB

Creation of a Chargeback Job is done through the Chargeback Job Wizard. After completing the creation wizard the job will be sent to the target NSS Service to be executed according to its schedule.

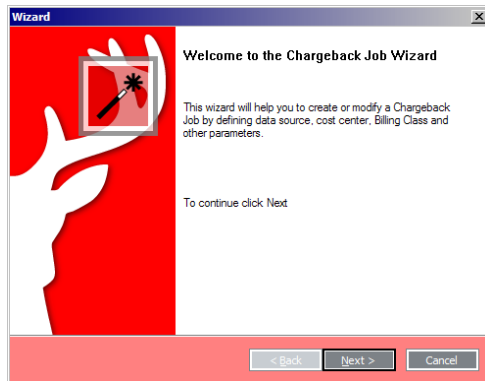
Chargeback Jobs require one or more Billing Class to convert usage data into cost. A default Billing Class is created automatically. Additional Billing Classes can be added, either when

defining the Chargeback Job or in the Billing Class Configurator as describe in “Billing classes” on page 28

To define a Chargeback Job:

**7** On the Jobs page, click the Add button. The Chargeback Job Wizard appears.

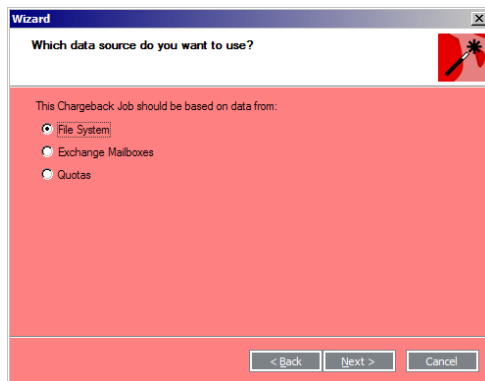
Figure 11: The Chargeback Job Wizard



**8** On the data source page, you can select which kind of usage data this Job will collect. Click Next

If you choose Quotas, make sure that Northern Quota Server is installed on the target Host and that you have a license for both Storage Chargeback and Quota Server for that Host. If you choose MS Exchange mailboxes make sure that you have created an Exchange profile as described in “Configuring for MS Exchange Jobs” on page 16.

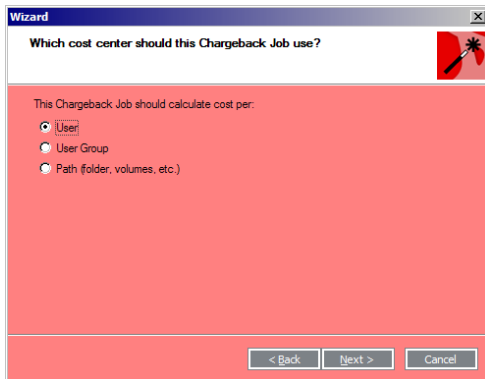
Figure 12: The data source page



- 9 On the Cost Center page, you can select how to associate the data source usage data to a given Cost Center.

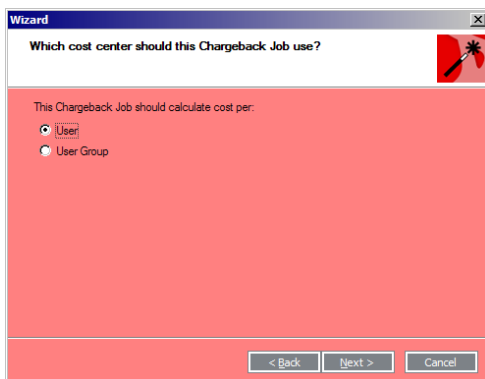
File System and Quotas Jobs provide the same Cost Center, i.e. User, User Group or Path.

Figure 13: Cost Center page for File System and Quota Jobs



For MS Exchange Jobs the Cost Centers are limited to User and User Group.

Figure 14: Cost Center page for MS Exchange Jobs




User Group always refers to the primary group. In other words, when selecting User Group, Storage Chargeback will get the usage information for all objects in the data source, retrieve the owner of the objects and resolve its primary group. The usage data is then summed up by primary group.

- 10 On the Path page you can select where the Job will retrieve data from and where the Job will be stored.

The top drop-down list presents you with all the available Hosts for the Job. A Host is available if:

- It's currently running the NSS Service
- It's has been added through the Host Management Configurator (See "Adding a Host" on page 13)

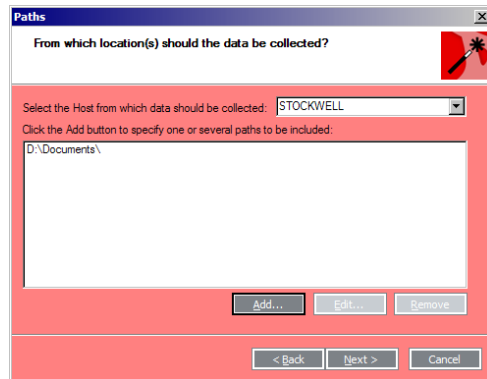
The Host that you select in this drop-down list will be the Host that will execute the Chargeback Job.

 *Notice! You cannot change the target Host after completing the Wizard. Since the Job definition will be added to the NSS Service on that Host, the only way to "move" it is to recreate a similar Job on another Host and delete the original Job.*

For File System Jobs, you can add several paths, even on multiple hosts. The host where the job is defined must have Northern Storage Suite installed, other hosts may be scanned remotely.

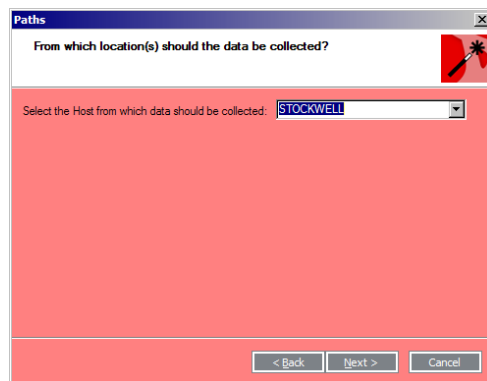
You can add, edit and remove a path with the buttons below. Adding a path brings up a Path browser. This browser will actually talk to the NSS Service on the selected Host to be able to show you all paths available to the service [since some paths might be accessible from the service account but not from the account running the GUI].

Figure 15: Path Selection page for File System Jobs



For Quotas and MS Exchange Jobs, you can only choose one Host, not individual path on that Hosts. Those Jobs will incorporate all available data (i.e. all Quotas or all mailboxes) from that Host regardless of path.

Figure 16: Server Selection for Quotas and MS-Exchange Jobs



**11** On the Billing Class page you assign a Billing Class for each path of the Job.

To assign a Billing Class for File System Jobs:

- select the path to modify
- select the Billing Class from the drop-down list

The path is now bound to that Billing Class

You can also define a new Billing Class from this page by clicking the "New..." button. Refer to "Adding a Billing Class" on page 30 for a description of the Billing Class Wizard. After the new Billing Class has been created, it will be available as a new choice in the drop-down list.


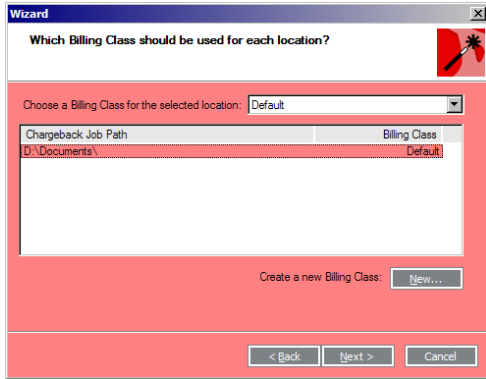
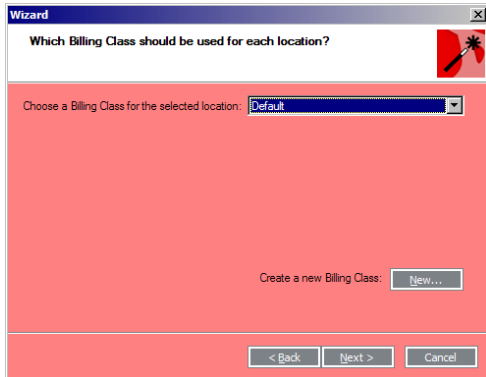
 **Notice!** Billing Classes created from the Job Wizard are global Billing Classes. They are then visible to all Jobs. In any aspect, they are the same as the Billing Classes created from the Billing Class Configurator.

Figure 17: Billing Class Page for File System Jobs



For Quotas and MS Exchange Job, since the path is a unique server, you just need to pick one Billing Class from the drop-down list.

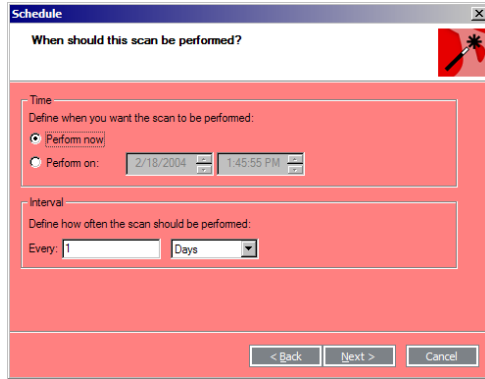
Figure 18: Billing Class Page for Quotas and MS Exchange Job




**12** On the Scheduling page, you select when the Job should be run and how it should be repeated.

If you select "Generate Now", the job will be executed as soon as it is added to the Host. Make sure that the repetition time leaves enough time for the Job to finish. The scheduling engine in Storage Chargeback will make sure that two Jobs cannot be executed at the same time [to avoid I/O collapse].

Figure 19: Scheduling page



 **Notice!** Under Windows 2003, the messenger service is disabled by default. This means that Popup notifications will fail if the NSS service or the recipient are using a Windows 2003 machine. The messenger service needs to be enabled manually through the Services Administrative Tool.

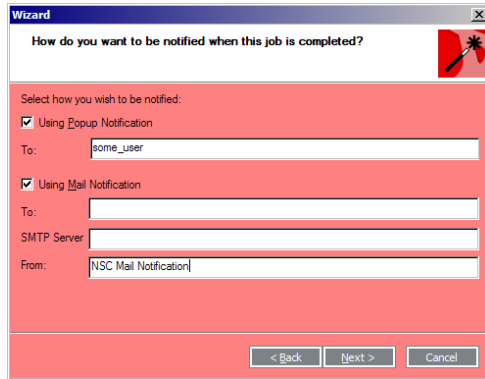
**13** On the notification page, you enter who should be notified.

You can send both a Popup and an email for the same Job. The notification is sent after the Job is completed with a status of the execution [succeeded or failed].

Popup can be send to either a user name or a computer name.

Mail notification can be sent to any valid email address reachable from the selected SMTP server. Make sure that you enter a valid user name or email address in the "From" field.

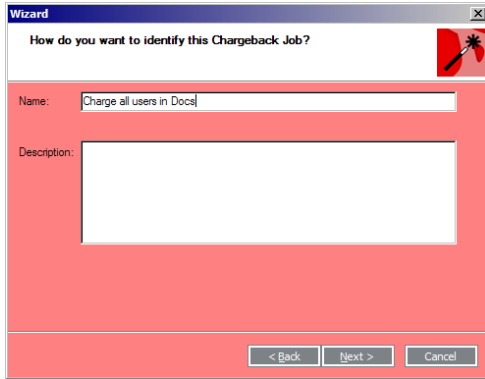
Figure 20: Notification Page



**14** On the Identification page, you enter the Job Name and Description

Only the Name is mandatory. The Description is used as memo information and does not need to be entered [or it can be entered when the Job is edited].

Figure 21: Identification Page



**15** You have reached the Finish page. Clicking the Finish button will effectively add the Job to the target Host.

### REMOVING A CHARGEBACK JOB

From the Jobs page you can remove a Chargeback Job by selecting it in the list and clicking the "Remove" button.

Removing the Chargeback Job will only remove the Job definition from the target Host. The Job results and their data will stay in the database. Also, Billing Classes defined during the creation of the Job are not affected by the removal.

### EDITING A CHARGEBACK JOB

From the Jobs page, you can edit a Chargeback Job by either:

- selecting it in the list and clicking the "Edit..." button
- or double clicking the Job entry in the list

Editing a Job will start the Job Wizard again. All parameters can be changed when editing except the target Host and the name. So a Chargeback Job that was defined to collect data from a Host A cannot be changed to collect data from another Host B. To change the location of a Chargeback Job you need to create a new Job.



*Notice! The name cannot be changed after creation of the Chargeback Job. Also, even if Storage Chargeback will accept Jobs with the identical name, it is strongly recommended, for clarity, to use unique names.*

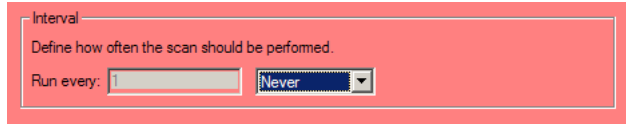


*Notice! You cannot undo the removal of Chargeback Job. Also if you delete a Job and add a new Job with the same parameters, Storage Chargeback will still consider the two Jobs as different. Even if you have used the same Name, the two Chargeback Jobs will receive different internal IDs.*

## Disabling a Chargeback Job

To disable a Chargeback Job just edit the Job as described above and choose the "Never" option in the scheduling page.

Figure 22: Never option in the Scheduling page



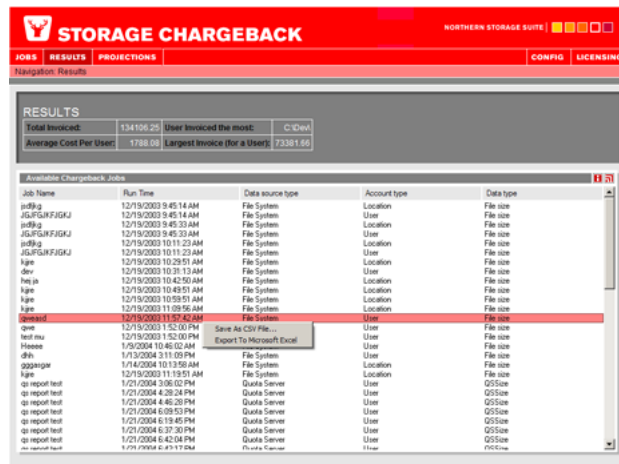
Interval  
Define how often the scan should be performed.  
Run every:  **Never**

## VIEWING A CHARGEBACK JOB RESULT

When a Job has been defined, it will start executing according to its scheduling details. Every time the Job is run it will create one Job Result.

All operations for a Chargeback Job Result are done from the Results view.

Figure 23: Storage Chargeback Results Page



**STORAGE CHARGEBACK** NORTHERN STORAGE SUITE

Navigation: Results

**RESULTS**

Total Invoiced:	134196.25	User Invoiced the most:	CIDevA
Average Cost Per User:	1788.00	Largest Invoice (for a User):	73381.66

Available Chargeback Jobs

Job Name	Run Time	Data source type	Account type	Data type
jd9lg	12/19/2003 9:45:14 AM	File System	Location	File size
JGwFwFJGKJ	12/19/2003 9:45:14 AM	File System	User	File size
jd9lg	12/19/2003 9:45:33 AM	File System	Location	File size
JGwFwFJGKJ	12/19/2003 9:45:33 AM	File System	User	File size
jd9lg	12/19/2003 10:11:23 AM	File System	Location	File size
JGwFwFJGKJ	12/19/2003 10:11:23 AM	File System	User	File size
kgie	12/19/2003 10:28:51 AM	File System	Location	File size
dev	12/19/2003 10:31:13 AM	File System	User	File size
helsh	12/19/2003 10:42:50 AM	File System	Location	File size
kgie	12/19/2003 10:49:51 AM	File System	Location	File size
kgie	12/19/2003 11:08:56 AM	File System	Location	File size
demard	12/19/2003 11:07:42 AM	File System	User	File size
qee	12/19/2003 1:52:00 PM	Save As CSV File...	User	File size
test.mta	12/19/2003 1:52:00 PM	Export to Microsoft Excel	User	File size
Hesee	1/9/2004 10:46:02 AM	File System	User	File size
qhu	1/13/2004 3:11:00 PM	File System	Location	File size
qppqppw	1/14/2004 10:13:58 AM	File System	Location	File size
kgie	12/19/2003 11:18:51 AM	File System	Location	File size
qi report test	1/21/2004 3:06:02 PM	Quota Server	User	QSSize
qi report test	1/21/2004 4:28:28 PM	Quota Server	User	QSSize
qi report test	1/21/2004 4:46:20 PM	Quota Server	User	QSSize
qi report test	1/21/2004 6:09:53 PM	Quota Server	User	QSSize
qi report test	1/21/2004 6:15:46 PM	Quota Server	User	QSSize
qi report test	1/21/2004 6:37:30 PM	Quota Server	User	QSSize
qi report test	1/21/2004 6:42:04 PM	Quota Server	User	QSSize
qi report test	1/21/2004 6:47:17 PM	Quota Server	User	QSSize

From this view you can:

- Review all Job Results
- Jump to the data of one specific Job Result.
- Export the results

## Reviewing Job Results

The Results view present a list of Job Results in the active database. By default, the Job Results are sorted by Run Time. You can change the sorting order by clicking on the heading of any column.

The list displays the following information:

ATTRIBUTE	DESCRIPTION
<b>Job Name</b>	This is the name that was entered during in the Chargeback Job definition
<b>Run Time</b>	Data and time of when the Chargeback Job finished executing.
<b>Data source Type</b>	Type of data collected by the Chargeback Job: <ul style="list-style-type: none"> <li>- File System</li> <li>- Quotas</li> <li>- MS Exchange</li> </ul>
<b>Cost Center</b>	This is the type of Cost Center used for the Chargeback Job: <ul style="list-style-type: none"> <li>- User</li> <li>- User Group</li> <li>- Location</li> </ul>
<b>Usage Data</b>	This is the usage data that was collected as the base of the calculated cost: <ul style="list-style-type: none"> <li>- File size</li> <li>- Quota size</li> <li>- Mailbox size</li> </ul>

### Viewing one specific Job Result

By double clicking a Job Result you can view the data of that specific Result. The Job Data page shows all the usage and cost data from a Job Result.


This data page will be slightly different according to what options have been selected during the definition of the Job. However it will always present the same conceptual information

ATTRIBUTE	DESCRIPTION
<b>Server</b>	This is the name of the server that executed the Chargeback Job.
<b>Path or Host</b>	The location of the Cost Center's storage (files, quotas or mailboxes). The heading is Path for all Chargeback Jobs using the file system data source. Otherwise the heading is Host.
<b>Billing Class</b>	The name of the Billing Class applied to the storage data of the Cost Center at this location.
<b>User or Cost Center</b>	The name of the Cost Center. The heading is User when the Chargeback Job was defined with User as the Cost Center. When the Cost Center type was user group or location, the heading is Cost Center and the content of this column is the same as the Path or Host column.
<b>Usage Data</b>	This is the usage data that was collected as base of the calculated cost: <ul style="list-style-type: none"> <li>- File size</li> <li>- Quota size</li> <li>- Mailbox size</li> </ul>

## Exporting the Job Results and Data

In both the Job Results and Job Data list you can right click anywhere to bring up the Exporting menu.

Figure 24: The Exporting menu of the Job Result and Data list



Both options will take all selected entries and copy their content to the proper output format. Note that the list is copied as it is, so sorting and column ordering are respected by the export.

## TRACKING CHARGEBACK JOBS

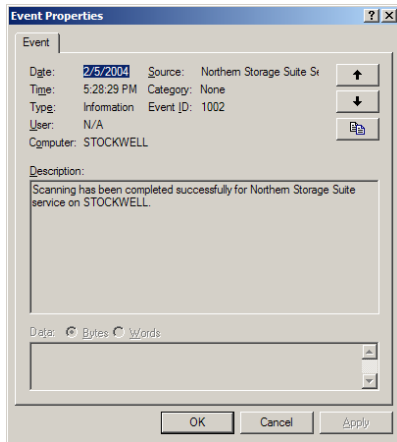
Storage Chargeback reports Job execution status in two ways:

- User notification
- Event Log

User Notification is configured during the definition of the Chargeback Job. Notifications are sent when the Job has finished executing and will report errors encountered during the execution of the Job, if any.


The Event Log will include a message signalling the end of the execution of the Job. Note that the message in the Event log will only report an error if the job completely failed. If the Job encountered errors but was able to finish, the Event Log message will describe a successful execution. Complete failure could occur, for instance, when no valid license is available.

Figure 25: Event Log message for successful Chargeback Job



## BILLING CLASSES

In Northern Storage Chargeback you can define a variety of different pricing or cost patterns, called Billing Classes. The Billing Class can be used to define multi-tiered pricing models. The Billing Classes can be given descriptive names and you can assign the Billing Classes arbitrarily to different storage points [volume by volume or per path, quota or equivalent].


 Notice! To be able to Export to Microsoft Excel, you will need to have it installed on your client machine (i.e. on the machine where you are running the Storage Chargeback GUI).

## ABOUT MULTI-TIERED PRICING MODELS

When you create a Billing Class, you define a cost formula. In the most simple case, you define a completely linear formula. That is  $x$  megabytes costs  $y$  credits, irrespective of how much storage is used. Or, you define a formula that states that: "up till 250 megabytes you pay one credit per megabyte, but after 250 megabytes you pay half a credit for every megabytes", etc. You can have any number of such brackets.

Example: in the illustration below you see a Billing Class with three brackets [or tiers]. In the first usage bracket [T1] you pay your usage in megabytes times 1.67 credits. In the second bracket [T2] you pay 0.75 credits for any megabytes used above the T1 threshold. In the third bracket you pay 0.25 credits for everything above the T2 threshold. It is important that the already accrued cost stacks with any cost attained in a higher bracket. The cost of the usage of all brackets contribute to the total, summary cost of a particular amount of storage use.

Figure 26: Billing Classes and multiple thresholds



Another aspect of the Billing Classes is, of course, that you can assign different Billing Classes to different storage facilities. You can even have multiple Billing Classes on the same drive. This is made possible since you can assign different Billing Classes to different paths.

Typically, you would use Billing Classes to mirror the costs you have for supplying different kinds of storage media. You can easily create a Billing Class for your standard storage, and one for premium storage. Premium storage could then be either more expensive in terms of the hardware itself being more expensive, or as a result of higher costs for maintenance resulting from more ambitious backup procedure, or higher availability standards [clustering, mirroring, etc.].

## ABOUT THE DEFAULT BILLING CLASS

Northern Storage Chargeback has a Default Billing Class. This Billing Class has the following cost scheme:

- If the usage is between zero [0] megabytes and 50 megabytes the per megabyte price is zero [0].
- If the usage is between 50 megabytes and 250 megabytes the per megabyte price is 1 credit
- If the usage is above 250 megabytes per megabyte price is 2 credits.

The cost scheme could also be expressed as:

- [0 - 50MB] 0 credits
- [50 - 250] 1 credit/MB
- [250 - infinity] 2 credits/MB

As any other Billing Class, the default Billing Class can be edited or even removed.

## ADDING A BILLING CLASS


Billing Classes are managed from either the Billing Class Configurator or the Chargeback Job Wizard. Both ways provide access to the same interface.

Billing Classes are stored into the database, and are thus global to all Hosts.

To define a Billing Class:

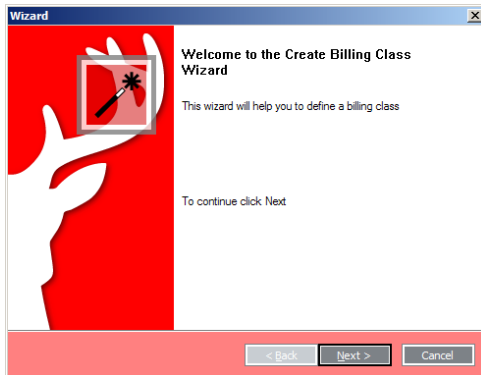
- 1 On the Config view click the Billing Class Management heading. The Billing Class Configurator opens

Figure 27: The Billing Class Configurator



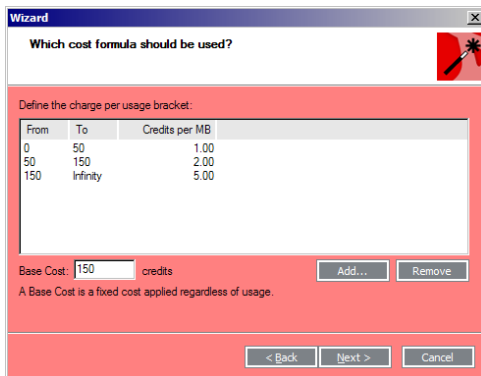
2 Click the Add button. The Billing Class Wizard appears.

Figure 28: The Billing Class Wizard




3 On the Bracket page you can define the threshold cost of the Billing Class

Figure 29: Bracket list page



A Billing Class is made of a base cost and zero or more price brackets. The base cost is entered in the "Base Cost" field. It is applied to each Cost Center regardless of its usage.

Figure 30: The Base Cost input field



Each bracket defines a cost multiplier for the data usage. To create a new bracket click the "Add..." button. This brings up the bracket definition dialog.

Figure 31: Bracket Definition Dialog



Notice! You need to use the dot sign to write fractional cost.

In this dialog you just need to enter the starting point of the Bracket and the cost multiplier. The Billing Class Wizard will automatically insert the new Bracket at its proper place, so that illogical definitions are impossible.

Below is the Bracket list as it appears in a new Billing Class after a new Bracket has been created. The Bracket below defines a cost of 1 credit per megabyte for any usage [0 to infinity].

Figure 32: A Billing Class with one Bracket

Define the charge per usage bracket:		
From	To	Credits per MB
0	Infinity	1.00

If you add a new Bracket to the Billing Class starting at 150 MB to the Billing Class with a cost of 5 credits per MB, the list will look like this:

Figure 33: A Billing Class with two Brackets

Define the charge per usage bracket:		
From	To	Credits per MB
0	150	1.00
150	Infinity	5.00

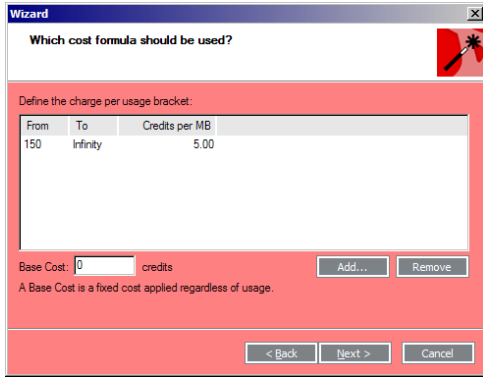
Now, if you add a new Bracket starting at 50 MB with a cost of 2 credits per MB, the Billing Class Wizard will automatically insert it at the proper place to create a new cost ratio between 50 and 150 MB, as below:

Figure 34: A Billing Class with three Brackets

Define the charge per usage bracket:		
From	To	Credits per MB
0	50	1.00
50	150	2.00
150	Infinity	5.00

Brackets do not have to start at 0 MB. The following Billing Class has only one Bracket and effectively allow the use of 150 MB of free storage.

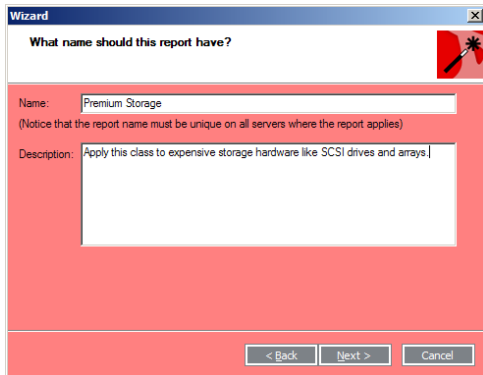
Figure 35: A Billing Class with free start storage



4 After creating the Billing Class formula you can go to the identity page. On this page you can enter the name and a description for the Billing Class.

Only the Name is mandatory on that page. The Description is used as memo information and does not need to be entered [or it can be entered later, when the Billing Class is edited]. Note that Billing Class names do not need to be unique, although using different classes with the same name might be confusing, to say the least.

Figure 36: Identity Page



## REMOVING A BILLING CLASS

Billing Classes cannot be removed. They can just be hidden. This is due to the fact that old Job Results might still depend on the Billing Class. If its definition was removed from the database, the Job data could not be presented correctly.

In the Billing Class Configurator you can click the "Remove" button to hide a Billing Class. It's now removed from the list but remains in the database.

For all purposes, except history data, the removed Billing Class has no effect on existing or new Billing Classes. For example, you can create a new Billing Class with the same name as the one you removed.

To permanently delete a Billing Class, you need to erase the data from the database. This operation is not recommended. It will disrupt all Job data and will save little space [a Billing Class definition uses less than 100 bytes of space in the database].


## EDITING A BILLING CLASS

In the Billing Class Configurator, you can edit a Billing Class by either clicking the "Edit" button or double clicking on the class in the list.

When editing a Billing Class you can change any of its parameters including its name.

If you change the Billing Class name, Jobs that refer to that Billing Class will be updated correctly, including Job Results generated in the past.

If you change the Billing Class brackets, old Chargeback Jobs that refer to that Billing Class will be unaffected. In other words, they will keep their calculated cost as it was previously defined.

 *Notice! Changing the brackets of a Billing Class will affect old and new Projections that uses that Billing Class.*

## PROJECTIONS

While Chargeback Jobs give you instant valuation for your storage usage, Northern Storage Chargeback projections offer in-depth analysis of cost trends.

### ABOUT PROJECTIONS

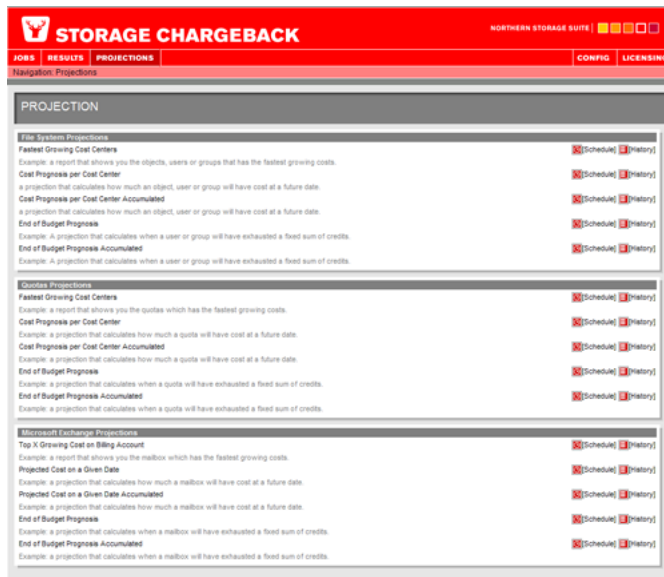
The Projections differ from the Chargeback Jobs in the respect that they do not generate any access to the current storage data. They will not scan the file system, query the quotas or enumerate MS Exchange mailboxes.

Instead, the Projections rely on the results of previously run Chargeback Jobs to report their data. This means that, to get meaningful, results you need to have run a several Chargeback Jobs.

However, the Projections do have the concepts of data source and Cost Center, but those are used to filter out the results Chargeback Jobs to take into account only those that have the same data source and Cost Center as the Projection. This is necessary to not mix-up different kinds of results [the cost of a user's file system usage cannot be directly compared with the cost of the users's mailbox storage].

Northern Storage Chargeback provides 5 main types of Projections that mine the Chargeback Job results in a different way. Each Projection type can then be applied to the three different data sources [File System, Quotas or MS Exchange]. The full list of Projection types is displayed on the main Projection view

Figure 37: The Projection View



The Projections are enterprise-wide reports. They can be applied to results from several Host. when defining a Projection you do not have to choose which Host the Projection will get the data from. The Projection always get its data from the database.

Also, due to their enterprise-level nature, all Projection definitions are stored and executed on the Master Server.

All Projection types are managed in the same way from the Projections page. The sections:

- “Adding a Projection” on page 35
  - “Removing a Projection” on page 39
  - “Editing a Projection Definition” on page 39
  - “Viewing a Projection Result” on page 40
- will describe how to perform the basic administration tasks regardless of the Projection type.

To get a detailed description of the behavior of each Projection type refer to:

- “Cost Prognosis per Cost Center” on page 46
- “Cost Prognosis per Cost Center (Accumulated)” on page 49
- “End-of-Budget Prognosis” on page 47
- “End-of-Budget Prognosis (Accumulated)” on page 51

## ADDING A PROJECTION

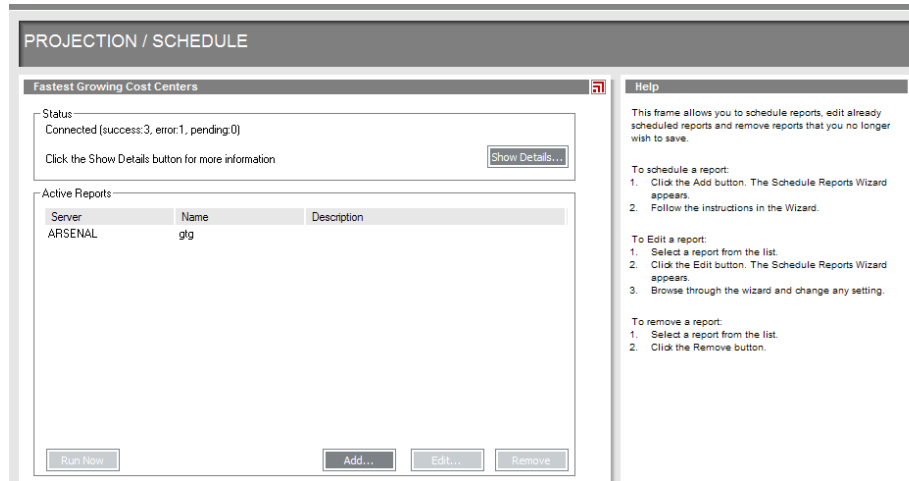
Creation of a Projection is done from the Schedule page of that Projection type. After completing the wizard, the Projection definition will be sent to the Master Server.

Projections require prior-results from one or more Chargeback Job of the same data source and Cost Center as described above.

To define a Projection:

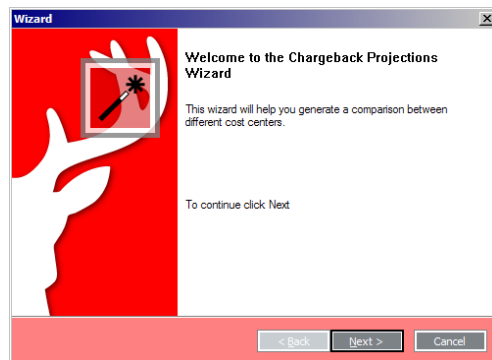
- 1 On the Projections page click the "Schedule" link of the chosen Projection type. Each Projection type has an individual "Schedule" and "History" link. Clicking on any of those links will bring you to the Schedule or History page specific to that Projection type.

Figure 38: The Schedule Page for a given Projection type



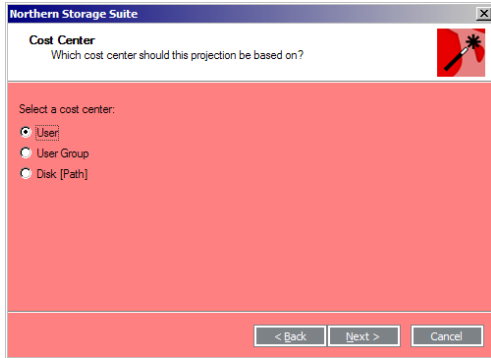
- 2 On the Projection Schedule page click the "Add..." button to bring up the Chargeback Projections Wizard. Click "Next" on the welcome page.

Figure 39: The Chargeback Projections Wizard



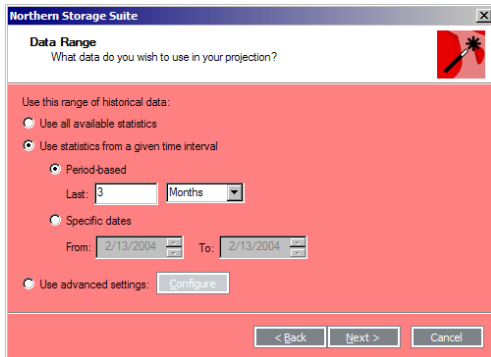
- 3 On the Cost Center page, you select which object or person to charge. This is actually a filter, since it will make the Projection only consider Chargeback Jobs with the same Cost Center.

Figure 40: The Cost Center page



- 4 On the Properties page, you choose attributes specific to the current Projection type. Each Projection type has a different attribute page. The current section focuses on general operations for any Projection, the pages are described in detail in "Using Projections" on page 44.
- 5 The Data Range page allows you to specify which Chargeback Job results you want to consider for this Projection. As mentioned before, the Projection will only take into account Chargeback Job results that have been generated with the same data source and Cost Center. The Data Range is an additional settings to filter out Chargeback Job results based on their running date.

Figure 41: The Data Range page



You can choose the Data Range in three ways:

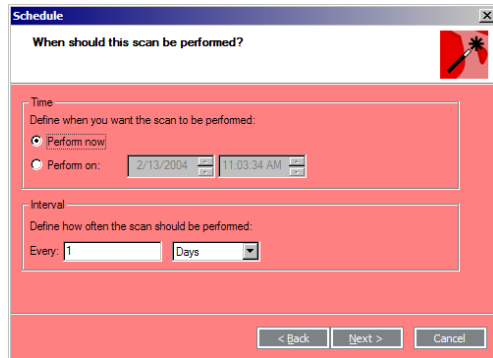
- Using all data
- Using static (fixed) dates
- Using dynamic (variable) dates

Northern Storage Chargeback provides very powerful ways to specifying which data you want to use for your Projection. Refer to "Defining a Data Range" on page 41 for a complete description of those options.

- 6 On the Scheduling page, you select when the Projection should be run and how it should be repeated.

If you select "Perform Now", the Projection will be executed as soon as it is added to the Master Server. Make sure that the repetition time leaves enough time for the Projection to finish. The scheduling engine in Storage Chargeback will make sure that two Projections cannot be executed at the same time [to avoid I/O collapse].

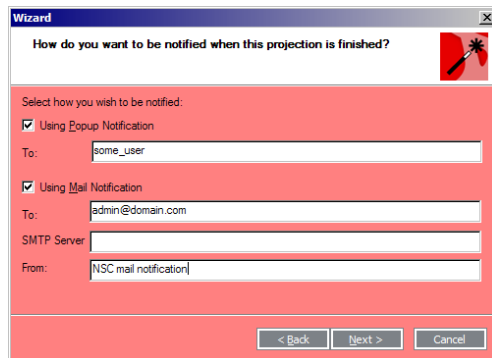
Figure 42: Scheduling page



- 7 On the notification page, you enter who should be notified.

You can send both a Popup and an email for the same Projection. The notification is sent after the Projection is completed with a status of the execution [succeeded or failed].  
Popup can be sent to both a user name or a computer name.  
Mail notification can be sent to any valid email address reachable from the given SMTP server.  
Make sure that you enter a valid user name or email address in the "From" field.

Figure 43: Notification Page

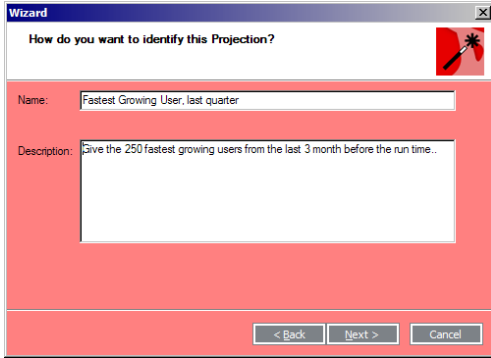


*Notice! Under Windows 2003, the messenger service is disabled by default. This means that Popup notification will fail if the NSS service or the recipient are using a Windows 2003 machine with default configuration. The messenger service needs to be enabled manually through the Services Administrative Tool.*

8 On the Identification page, you enter the Projection Name and Description

Only the Name is mandatory on that page. The Description is used as memo information and does not need to be entered [or it can be entered when the Projection is edited].

Figure 44: Identification Page



9 You have reached the Finish page. Clicking the Finish button will effectively add the Projection to the Master Server.

### REMOVING A PROJECTION

From the Projection's Schedule page you can remove a Projection by selecting it in the list and clicking the "Remove" button.

Removing a Projection will only remove the Projection definition from the Master Server. The Projection results and its data will stay in the database. Chargeback Jobs result used by the Projection are not affected by the removal.

### EDITING A PROJECTION DEFINITION

From the Projection's Schedule page you can edit a Projection definition by either:

- Selection it in the list and clicking the "Edit..." button
- Double clicking it in the list

Editing a Projection will start the Chargeback Projection Wizard again. All parameters can be changed during the editing except its name.


### Disabling a Projection

To disable a Projection just edit its definition as described above and choose the "Never" option in the scheduling page.



*Notice! You cannot undo the removal of a Projection. Also if you delete a Projection and add a new Projection with the same parameters, Storage Chargeback will still consider the two Projections as different. Even if you have used the same Name, the two Projections will have different internal IDs.*

Figure 45: Never option in the Scheduling page

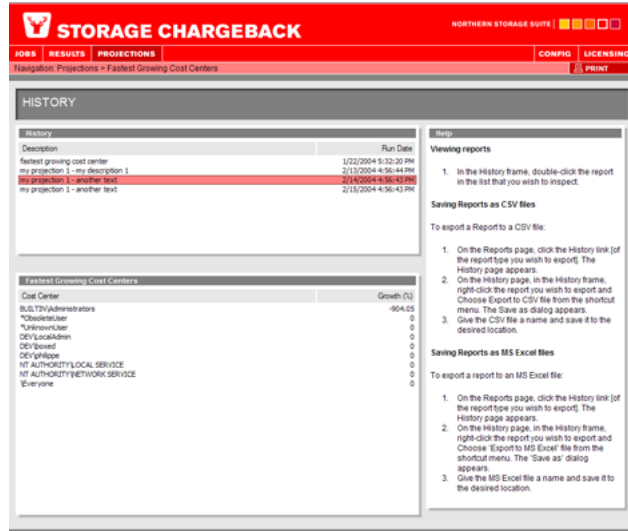


## VIEWING A PROJECTION RESULT

When a Projection has been defined, it will start executing according to its scheduling details. Everytime the Projection is run it will create one result.

All operations for a Projection result are done from the Projection's History page

Figure 46: A Projection History Page



From this page, you can:

- Review all results of a Projection type
- Review all data of one specific Projection result
- Export the results or the data

## Viewing Projection Results

The top list of the Projection History page contains one entry per result for the Projection type. Note that contrary to Chargeback Jobs, you cannot see all the results from different Projections at the same time. The results, like the definition, are grouped by Projection type.

By default the results are sorted by Run Date. You can change the sorting order by clicking on the heading of any column.

The Projection result list presents the following attributes for each entry in the list:

Figure 47: Projection Results Attributes

ATTRIBUTE	DESCRIPTION
<b>Description</b>	This is the description that was entered during in the projection definition
<b>Run Date</b>	Data and time of when the execution of the Projection finished.

## Viewing a Projections data

The bottom list of the Projection History page shows all the data calculated during the processing of the Projection. It will show one entry per Cost Center applicable to the Projection.

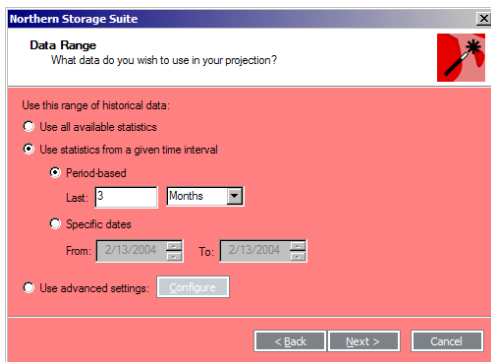
By default the list is unsorted. Actually the data is displayed as it was added to the database. First entry added to the database will be shown at the top of the list. You can change the sorting order by clicking on any column heading.

The Projection data list presents different attributes according to the Projection type. To get a full description of all the attributes, refer to the corresponding Usage section under "Using Projections" on page 44.

## DEFINING A DATA RANGE

One of the most important settings of a Projection is to specify which Chargeback Job results will be taken into account when calculating the projected values. This is done through the Data Range page of the Chargeback Projection Wizard.

Figure 48: The Data Range page



The Data Range is basically a Start and an End Date. However the way Storage Chargeback calculates those dates is exactly what the Data Range page allows use to configure.


You can specify the Data Range in three ways:

- Using all data
- Using basic settings
- Using advanced settings

### Full Data Range

You can choose to use all data by checking the "Use all available statistics" box.

Figure 49: The "Use all available statistics" option



This option will instruct the Projection to make no filtering based on the Chargeback Job results run time. All results with the proper data source and Cost Center will be considered.

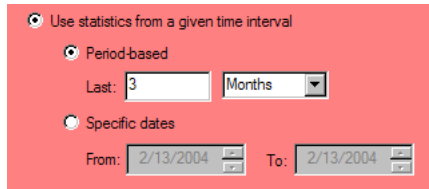


*Notice! Start date is inclusive but end date is exclusive. For example, a Projection using results from January 1st 2004 to January 1st 2005, will use the data collected during all of the 2004 year including the 1st of January to the 31st of December 2004..*

## Basic Data Range

Using static dates, you define a start and an end date.

Figure 50: The "Use statistics from a given time interval" option



Any Chargeback Job results that has been executed between this start and end date will be used by the Projection. All other results are ignored.

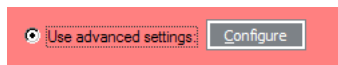
By choosing "Period-based", the end date is automatically set to the creation time of the Projection. The start date is set according to the period setting. For example, if the Projection is created on January 11th, 2004 and the period is set to 3 months, then all Chargeback Jobs results generated between November 11th, 2003 [that's 3 months before] and January 11th, 2004 will be included in the Projection. All other results will be ignored.

## Advanced Data Range

Northern Storage Chargeback supports truly dynamic Data Ranges. For example, you can define a Projection that will consider results from the first half of the year, regardless when the Projection is run. Or you could define a Projection that uses results from the last quarter, again regardless of when the Projection is executed.

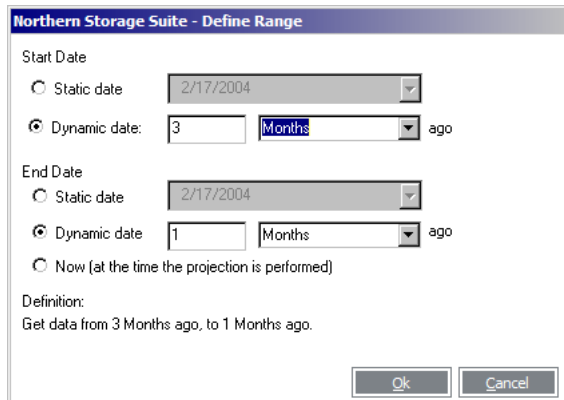
These settings are done through the "Use Advanced settings" option which, when selected enables the "Configure" button.

Figure 51: The "Use advanced settings"



Clicking the button brings up the Range Definition dialog that allows you to define the most advanced time ranges.

Figure 52: The Range Definition Dialog



*Notice! If a calculated date does not exist (like 30th of February), Northern Storage Chargeback will use the next available date (1st of March)..*



*Notice! There is a difference between "Last 1 Month" and "Last 30 days". Indeed, for an execution date of January 11th 2004, "Last 1 Month" will be from December 11th 2003 [i.e. the same day but of the previous month]. Whereas "Last 30 days" will be from December 12th 2003 [i.e. 30 running days before the 11th of January].*

This is done by telling Northern Storage Chargeback how to calculate the Start and End date of the range.

In the Range Definition Dialog you can set both Start Date and End Date to a fixed date. This option is the same as described above in the basic settings and is reproduced in the Range Definition Dialog for convenience.

The true power of this dialog comes when using the second set of choices: Dynamic Date. Checking the "Dynamic Date" option allows the user to enter a period count and length, for example 3 months.

When used for the Start Date, the period means the beginning of the last full period. So for a 3 month period, this will be the beginning of the previous 3 months. If the Projection is run on the 17th of June, the Data Range will start on the 1st of March.

When used for the End Date, the period means the end of the last full period. So for a 3 month period, this will be the end of the previous 3 months. If the Projection is run on the 17th of June, the Data Range will end on the 31st of May.

The table below shows how Storage Chargeback will interpret various typical settings.

Figure 53: Dynamic Data Range Examples

START DATE	END DATE	DATA RANGE
The examples below assume a first execution date of Wednesday 17th of June 2004.		
LAST 1 WEEK	LAST 1 WEEK	Monday, June 7th 2004 to Sunday, June 13th 2004. Everyday of last week.
LAST 2 WEEKS	LAST 1 WEEK	Monday, May 31st 2004 to Sunday, June 13th 2004. Everyday of the last fortnight.
LAST 4 WEEKS	LAST 1 WEEK	Monday, May 17th 2004 to Sunday, June 13th 2004. Everyday of the last 4 weeks.
LAST 1 MONTH	LAST 1 MONTH	Sunday, May 1st 2004 to Monday, May 31st 2004 Everyday of the last month.
The examples below assume a first execution date of 15th of December 2004.		
LAST 2 MONTHS	LAST 1 MONTH	October 1st 2004 to November 30th 2004. Everyday of the last two months.
LAST 2 MONTH	LAST 2 MONTH	October 1st 2004 to October 31st 2004. Everyday of the month before last month.
LAST 3 MONTHS	LAST 2 MONTHS	September 1st 2004 to October 31st 2004. Each day within a two month period, ending two months into the past.
LAST 1 QUARTER	LAST 1 QUARTER	July 1st 2004 to September 30th 2004. Everyday of the last calendar quarter.



*Notice! Northern Storage Chargeback will use your regional settings to understand what is the beginning and end of the week. So that US settings, the week is from Sunday to Saturday.*

START DATE	END DATE	DATA RANGE
LAST 3 MONTH	LAST 1 MONTH	September 1st 2004 to November 30th 2004. Everyday of the last 3 running months.
LAST 90 DAYS	LAST 1 DAY	September 15th 2004 to December 14th 2004. Everyday of the last 90 running days.
The examples below assume a first execution date of 10th of July 2005.		
LAST 1 YEAR	LAST 1 YEAR	January 1st 2004 to December 31st 2004. Everyday of last year.
LAST 4 QUARTERS	LAST 1 QUARTER	July 1st 2004 to June 30th 2005. Everyday of the last 4 running quarters.
LAST 1 HALF YEAR	LAST 1 HALF YEAR	July 1st 2004 to December 31st 2004. Everyday of the last Half year.
LAST 2 HALF YEAR	LAST 2 HALF YEAR	January 1st 2004 to June 30th 2004. Everyday of the first Half of last year.

## USING PROJECTIONS

As mentioned before Northern Storage Chargeback provides 5 types of Projections. Although, most settings and operations are similar for each type, some information is completely specific. The differences and specifics are described below.

The Projection uses the Data Range dates to retrieve usage data from appropriate Chargeback Jobs. The usage data is then extrapolated according to the formula described in "About Projection Extrapolation" on page 66. The extrapolated usage data is then passed through all required Billing Classes to calculate a final projected cost. One cost is returned for each Cost Center.

## FASTEST GROWING COST CENTERS

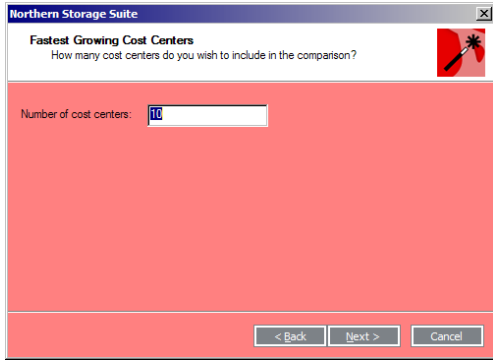
This Projection type presents a list of the Cost Centers that exhibit the highest increase of cost.

Contrary to the other Projection types, the Fastest Growing Projection does not return values [cost or date] for future usage. Instead it describes the current growth speed. The speed is expressed in percent. Still the Projection uses the Data Range dates to retrieve usage data from appropriate Chargeback Jobs. One speed is returned for each Cost Center.

The growth speed is expressed in credit per day. So a speed of 250 for a Cost Center A signifies that A is costing 250 more credit every month.

The parameters for a Fastest Growing Projection are defined in the Chargeback Projection Wizard on the property page.

Figure 54: The Properties Page for a Fastest Growing Projection



The "Number of Cost Centers" is used to set maximum limit on the number of Cost Centers to return in the list. Use -1 to specify that you want all Cost Centers.

When the Projection is executed its data will be displayed in lower list of the Projection History Page as described in "Viewing a Projection Result" on page 40.


 *Notice! Storage Chargeback needs to evaluate all Cost Centers to be able to calculate the result. So, the "Number of Cost Centers" is a result option. It will reduce the storage used by the result in the database and the size of the data in the GUI, but will not speed up the data collection*

Figure 55: Projection History Page for Fastest Growing

Fastest Growing Cost Centers	
Cost Center	Growth
c:\dev	12.4
c:\mir	-5.23

The page displays a list of all data with one entry per Cost Center with the following attributes:

Figure 56: Fastest Growing Attributes

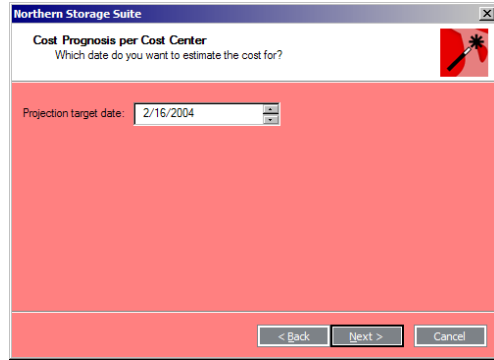
ATTRIBUTE	DESCRIPTION
COST CENTER	The name of the Cost Center the projected cost applies to. Either a user, a user group or a location depending of the settings of the projection
GROWTH	The calculated speed of increase or decrease. Positive growth speed indicate an increase in cost per month. Negative speed indicate a decrease or the cost. And null growth indicates that the change in cost for the Cost Center is constant [this does not reflect that actual cost, merely any change in cost].


## COST PROGNOSIS PER COST CENTER

This Projection type calculates how much a Cost Center will cost at a future date.

The parameters for a Cost Prognosis Projection are defined in the Chargeback Projection Wizard on the property page.

Figure 57: The Properties Page for a Cost Prognosis Projection



 *Notice! A Cost Prognosis Projection will never return a negative cost. Since Storage Chargeback uses exponential extrapolation, a decreasing cost trend will reach a value of zero at infinity.*

The "Projection Target Date" will be used as the extrapolation date, i.e. Northern Storage Chargeback will analyse the current cost pattern of each Cost Center to calculate their theoretical cost at the "Projection Target Date".

When the Projection is executed, its data will be displayed in the lower list of the Projection History Page as described in "Viewing a Projection Result" on page 40.

Figure 58: Projection History Page for Cost Prognosis

Cost Prognosis per Cost Center		
Cost Center	Cost	Projection date
c:\program files\LeechFTP	1.00	1/22/2007
c:\program files\MSDE	2.88	1/22/2007
C:\TEMP\	50.00	1/22/2007
c:\mir	52.86	1/22/2007
c:\program files\microsoft visual studio	892.94	1/22/2007
c:\dev	52175.17	1/22/2007
C:\Dev\	73381.66	1/22/2007

The page displays a list of all data with one entry per Cost Center with the following attributes:

Figure 59: Cost Prognosis Attributes

ATTRIBUTE	DESCRIPTION
COST CENTER	The name of the Cost Center the projected cost applies to. Either a user, a user group or a location depending of the settings of the projection
COST	The extrapolated projected cost in credits for a Cost Center.
PROJECTION DATE	The date entered in the "Projection Target Date" field of the Projection properties.

## END-OF-BUDGET PROGNOSIS

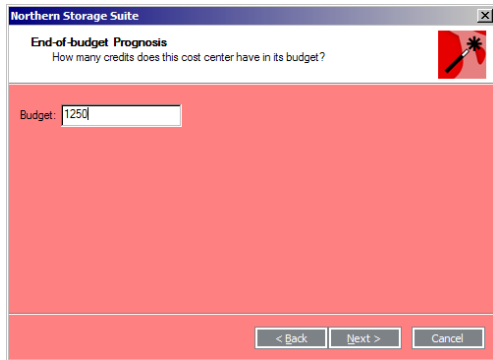
This Projection type calculates when a Cost Center will run out of budget. In a way, this is the reverse of the Cost Prognosis. Where the Cost Prognosis calculate a projected cost for a given date, the End-of-Budget returns a projected date for a given cost.

For example, a department is allocated a budget of 2000 credits for its storage usage. If that department is charged 250 credits a month for their usage, the Projection will show that it will run out of budget in 8 months.

This example has been simplified by ignoring any variation of the usage. Northern Storage Chargeback, on the other hand, is able to take into account the variation of the usage. Storage Chargeback will also take into account any Billing Classes used for the past Chargeback Jobs and will do so for each Cost Centers so that one results is returned for each Cost Center.

The parameters for a End-of-Budget Projection are defined in the Chargeback Projection Wizard on the property page.

Figure 60: The Properties Page for a End-of-Budget Projection



The "Budget" is simply the allocated value for each Cost Center.

When the Projection is executed its data will be displayed in the the lower list of the Projection History Page as described in "Viewing a Projection Result" on page 40.

Figure 61: Projection History Page for End-of-Budget Prognosis

End of Budget Prognosis		
Cost Center	Date (estimated)	Given Cost
C:\TEMP\	1/22/2004	800.00



Notice! It is possible for a Cost Center to never reach the budget. In that case, the "Date (estimated)" value will be empty..

The page displays a list of all data with one entry per Cost Center with the following attributes:

Figure 62: End of Budget Prognosis Attributes

ATTRIBUTE	DESCRIPTION
COST CENTER	The name of the Cost Center the projected date apply to. Either a user, a user group or a location depending of the settings of the projection
DATE (ESTIMATED)	The extrapolated projected date for a Cost Center.
GIVEN COST	The value entered in the "Budget" field of the Projection properties.

### ACCUMULATE PROJECTIONS

The above Projections give either the cost that a Cost Center will reach at a given date or the date when a Cost Center will reach a given cost. Those models might not represent two very popular billing schemes; utilities usage and renting.


When using utilities, like electricity, the user consumption is retrieved each period [e.g. month]. If the user is billed at the end of the year, his total bill will be will the accumulated cost of all monthly installements.

Renting is basically the same scheme except that the cost is constant from one period to another. Still the actual cost is the sum of the cost at each period.

Northern Storage Chargeback can reproduce those billing schemes using the accumulated Projections. This processing is represented on the graphs below.


In Figure 63, the past usage between the dates A and B is extrapolated into a future usage for each selected periods P1 to P7.

Figure 63: Extrapolated cost at given periods



The extrapolated cost are then added together to return a final accumulated cost at P7, as shown in the Figure 64.

Figure 64: Accumulated Sum of extrapolated costs



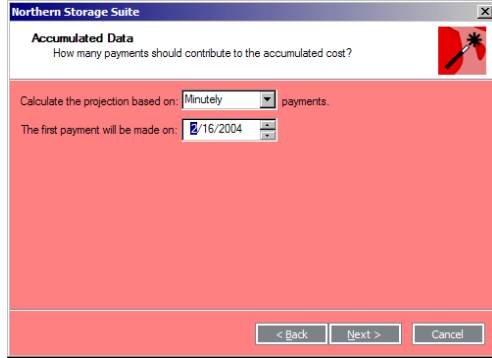
At P7, the accumulated is much higher than what would have been extrapolated with the non-accumulated Projections [shown by the dashed line].

### **COST PROGNOSIS PER COST CENTER (ACCUMULATED)**

Similar to the non Accumulated version, the Cost Prognosis per Cost Center calculate a [accumulated] cost for a given date.

The parameters for an accumulated Cost Prognosis Projection are the same as for the non accumulated version with the addition of the period definition

Figure 65: Period Definition Page for an Accumulated Projection



On this page you define the length of a period and the date of the first payment [which is the end of the first period].

When the Projection is executed its data will be displayed in the lower list of the Projection History Page as described in "Viewing a Projection Result" on page 40.


Figure 66: Projection History Page for accumulated Cost Prognosis

Cost Prognosis per Cost Center Accumulated		
Cost Center	Cost (Accumulated)	Given Date
c:\program files\LeechFTP	157.00	1/25/2007
c:\program files\MSDE	451.38	1/25/2007
c:\mir	8299.93	1/25/2007
C:\TEMP\	8508.95	1/25/2007
c:\program files\microsoft visual studio	140190.92	1/25/2007
c:\dev	8191501.85	1/25/2007
C:\Dev\	11520920.42	1/25/2007

The page displays a list of all data with one entry per Cost Center with the following attributes:

Figure 67: Accumulated Cost Prognosis Attributes

ATTRIBUTE	DESCRIPTION
COST CENTER	The name of the Cost Center the projected date apply to. Either a user, a user group or a location depending of the settings of the projection
COST (ACCUMULATED)	The extrapolated projected cost for a Cost Center at a given date.
GIVEN DATE	The value entered in the "Projection Target Date" field of the Projection properties.

 **Notice!** An accumulated Cost Prognosis Projection will never return a negative cost. Since Storage Chargeback uses exponential extrapolation, a decreasing cost trend will reach a value of zero at infinity.

## END-OF-BUDGET PROGNOSIS (ACCUMULATED)

End-of-Budget Prognosis can also be used with accumulation as for Cost Prognosis. As for the non-accumulated version, the End-of-Budget Projection calculate a date for a given cost [instead of a cost for a given date as in the Cost Prognosis].

The parameters for an accumulated End-of-Budget Projection are the same as for the non-accumulated version with the addition of the period definition

Figure 68: Period Definition Page for an Accumulated Projection

On this page you define the length of a period and the date of the first payment [which is the end of the first period].

When the Projection is executed its data will be displayed in the lower list of the Projection History Page as described in “Viewing a Projection Result” on page 40.


Figure 69: Projection History Page for accumulated End to Budget Prognosis

End of Budget Prognosis Accumulated		
Cost Center	Date (estimated)	Given Cost
C:\Dev\	1/22/2004	1000.00
C:\TEMP\	1/22/2004	1000.00
c:\dev	1/22/2004	1000.00
c:\program files\microsoft visual studio	1/29/2004	1000.00
c:\mir	5/27/2004	1000.00
c:\program files\MSDE	9/16/2010	1000.00
c:\program files\LeechFTP	3/23/2023	1000.00

The page displays a list of all data with one entry per Cost Center with the following attributes:

Figure 70: Accumulated End of Budget Prognosis Attributes

ATTRIBUTE	DESCRIPTION
COST CENTER	The name of the Cost Center the projected date apply to. Either a user, a user group or a location depending of the settings of the projection
DATE (ESTIMATED)	The extrapolated projected date when the Cost Center will reach the given cost.
GIVEN COST	The value entered in the "Budget" field of the Projection properties.

 *Notice! It's possible for a Cost Center to never reach the budget. In that case, the "Date (estimated)" value will be empty.*



# TECHNICAL

## SQL DATABASE STRUCTURE

Northern Storage Chargeback was designed to be open and extendable. All data stored in by Storage Chargeback in the SQL database is public and easily integrated with a third-party database engine, like an accounting system.

This section describes the database structure created by Northern Storage Chargeback as well as the SQL operations used by the NSS Core Service to manage this data.

### CONCEPTS

Chargeback Jobs and Chargeback Projections are stored in two different ways. Both, however, will use the common Billing Classes and data Path tables [as described below]. Except for those two common tables, Jobs and Projections share no tables.

When it comes to the database, Storage Chargeback has only 5 Projection types [as opposed to 15 as seen in the client]. Each of the Projection types has a field to specify its data source: file system, MS-Exchange or Quota. For the actual database structure there is no difference between a file system "Fastest Growing Cost Centers" and a MS-Exchange "Fastest Growing Cost Centers".

As mentioned previously in this manual, Jobs have a definition [i.e. its settings] and one or more runs [i.e. the actual resulting data]. Each time the Job is run one row is added to the definition table and zero or more rows are created in the two Data tables. The reason why the definition is re-written to the database is that it might have changed between two runs of the Job. The database engine needs the history of the Job definitions to be able to run the Projections. This scheme also applies to Projections with one table for the Projection definition and one table for the actual results.

### TABLE DEFINITIONS

Northern Storage Chargeback uses 2 base tables, 3 Job tables and 2 tables for each Projection type [10 Projection tables in total]. All the tables are only created on their first use, so when the first Chargeback Job is defined, the definition table is created, etc...

The 2 base tables are:

- CBPathMap which is used for mapping paths to Billing Classes
- CBBillingClasses, which is used to store Billing Classes definition

The Chargeback Jobs tables are:

- CBJobDefinition
- CBJobData
- CBJobDataResult

The Chargeback Projections tables are:

- CBxxxDef, stored the definition for the Projections of type xxx
- CBxxxRst, stored all the results of the Projection type xxx

The possible names for the Projections are:

- ReportTop for the Fastest Growing Cost Center Projection [so there will be a CBReportTopDEF and CBReportTopRST table]
- ReportProjCost for the Cost Prognosis per Cost Center Projection
- ReportProjCostAcc for the Cost Prognosis per Cost Center Accumulated Projection
- ReportProjDate for the End of Budget Prognosis Projection
- ReportProjDataAcc for the End of Budget Prognosis Accumulated Projection

### Table CBPathMap

This table holds the mapping between Billing Classes and the Paths they apply to.

COLUMN	TYPE	DESCRIPTION
<b>DataSrc</b>	TINY INT	The type of the Path where the data to charge on is: - 0 = File System source - 1 = Northern Quota Server's quotas - 2 = MS Exchange mailboxes
<b>ValueType</b>	TINYINT	The type of the quantity used as usage: - 0 usage is the size of data [e.g. size of the files] - 1 usage is the count of data [e.g. count of files] The value 1 is reserved for future use.
<b>Path</b>	NVARCHAR(255)	The path where the Billing Classes will apply. If the DataSrc is 1 or 2, the Path is only a server name, otherwise it's any valid file system path.
<b>BillingClass</b>	INT	The id of the BillingClass using this path. This id maps to the field Id of the CBBillingClass table.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBPathMap(DataSrc tinyint,ValueType tinyint,Path
    nvarchar(255), BillingClass int,
CONSTRAINT CBPathMapUPath PRIMARY KEY(DataRc,ValueType,Path));
```

There is a 1-to-N relationship between the Billing Classes and the Path. Also, you can define two Paths for the same server providing there are using different Data Source [for example, a MS Exchange Jobs and a Quota Job both on server A and with different Billing Classes].

### Table CBBillingClasses

This table holds the definition for a Billing Class.

COLUMN	TYPE	DESCRIPTION
<b>Id</b>	INT	The internal Id of the Billing Class
<b>BillingClass</b>	NVARCHAR(255)	The name provided when defining the Billing Class

COLUMN	TYPE	DESCRIPTION
<b>Description</b>	NVARCHAR(255)	The description provided when defining the Billing Class
<b>Low</b>	NUMERIC(19,0)	The minimum usage a Cost Center must reach to fall into this Billing Class. For example a value of 10MB will means that a user with 11MB of storage will have 1 MB counted using the Billing Class definition.
<b>CostBase</b>	MONEY	The base cost for the Billing Class, i.e. a cost always apply regardless of the usage.
<b>CostPerMB</b>	MONEY	Cost per MB of usage above the Low value.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBBillingClasses (Id INT, BillingClass
    NVARCHAR(255), Description NVARCHAR(255), Low
    numeric(19,0), CostBase money, CostPerMB money),
CONSTRAINT LowLimit PRIMARY KEY (Id, Low);
```

Since the name is not an index, several Billing Classes can have the same name.

### Table CBJobDefinition

This table contains one row per Chargeback Job runs. A Chargeback Job which is run several times will have one row per run. This way if the Job definition has changed between two runs, the database engine will still be able to interpret the Job data properly.,

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	NUMERIC(19,0)	The internal ID of the Chargeback Job. This ID is unique in the table so for each run a Job will get a different ID.
<b>Name</b>	NVARCHAR(255)	The name of the Chargeback Job as defined when the Job was created.
<b>DataSrc</b>	TINYINT	The type of the Path where the data to charge on is: - 0 = File System source - 1 = Northern Quota Server's quotas - 2 = MS Exchange mailboxes
<b>AccountType</b>	TINYINT	The type of the Cost Center for that Job: - 0 = Windows User Account - 1 = Windows Primary User Group - 2 = Path/Location - 3 = Autopath, i.e. all first-level subfolders of a given path [only apply to File System data sources].
<b>ValueType</b>	TINYINT	The type of the quantity used as usage: - 0 = usage is in size of data [e.g. size of the files] - 1 = usage is in count of data [e.g. count of files] The value 1 is reserved for future use.

<b>COLUMN</b>	<b>TYPE</b>	<b>DESCRIPTION</b>
<b>Runtime</b>	DATETIME	The date and time the Job was run.
<b>Definition</b>	VARBINARY (8000)	The internal Job definition used by the NSS Core Service in binary form. This field is reserved for future use.
<b>AccountCount</b>	INT	Number of unique Cost Centers in the Job result.
<b>CollectionTime</b>	NUMERIC(19,0)	Number of milliseconds used to scan the data source, before making any calculations.
<b>CalculationTime</b>	NUMERIC(19,0)	Number of milliseconds used to calculate the final costs after the data was collected This number will usually be very small or even null for smaller Jobs.
<b>TotalValue</b>	NUMERIC(19,0)	Accumulated usage for all Cost Centers in the Job result.
<b>MaxValue</b>	NUMERIC(19,0)	Maximum usage of all Cost Centers in the Job result.
<b>MaxValueOwner</b>	NVARCHAR(255)	Name of the Cost Center with the maximum usage. In case of equality only one name is returned.
<b>MedianValue</b>	NUMERIC(19,0)	Median usage of all Cost Centers in the Job result.
<b>AverageValue</b>	NUMERIC(19,0)	Average usage of all Cost Centers in the Job result.
<b>CountAAvgValue</b>	INT	Number of Cost Centers which usage is strictly above the average usage of the Job result.
<b>CountBAvgValue</b>	INT	Number of Cost Centers which usage is strictly below the average usage of the Job result.
<b>TotalCost</b>	MONEY	Accumulated cost for all Cost Centers in the Job result.
<b>MaxCost</b>	MONEY	Maximum cost of all Cost Centers in the Job result.
<b>MaxCostOwner</b>	NVARCHAR(255)	Name of the Cost Center with the maximum cost. In case of equality only one name is returned.
<b>MedianCost</b>	MONEY	Median cost of all Cost Centers in the Job result.
<b>AverageCost</b>	MONEY	Average cost of all Cost Centers in the Job result.
<b>CountAAvgCost</b>	INT	Number of Cost Centers which cost is strictly above the average cost of the Job result.
<b>CountBAvgCost</b>	INT	Number of Cost Centers which cost is strictly below the average cost of the Job result.
<b>STATUS</b>	TINYINT	The state of the Job run: - 0 = scan is in progress - 1 = scan done, calculation in progress - 2 = job completed See text below for more information.

The NSS Core Service creates this table using the SQL statement:

```

CREATE TABLE CBJobDefinition(ID numeric(19,0), Name nvarchar(255),
  DataSrc TINYINT, AccountType TINYINT, ValueType TINYINT,
  RunTime datetime, Definition varbinary(8000), AccountCount INT,
  CollectionTime numeric(19,0), CalculationTime numeric(19,0),
  TotalValue numeric(19,0), MaxValue numeric(19,0), MaxValueOwner
  nvarchar(255), MedianValue numeric(19,0), AverageValue
  numeric(19,0), CountAAvgValue int, CountBAvgValue int,
  TotalCost money, MaxCost money, MaxCostOwner nvarchar(255),
  MedianCost money, AverageCose money, CountAAvgCost int,
  CountBAvgCost int, Status tinyint,
CONSTRAINT IDIDX PRIMARY KEY(Id));

```

The row in the CBJobDefinition table is created as soon as the Job starts to execute. At that point, the Status field has a value of 0. When the scan is finished and all usage has been collected, the usage is stored in the CBJobData table. The calculation process then starts and the Status is updated to value 1. The usage data is translated to a cost data using the Billing Class definitions and the result is stored in the CBJobDataResult table. At the end of this process, the Status is set to 2.

The statistical data of the Job run (like AccountCount,...) is only populated when the Status reaches "job complete" [status=2]. It is possible to have Jobs that never reach the completion status if the service is stopped while the Job is running or if the Job encounters a critical error [like database connectivity problems,...].

### Table CBJobData

This table holds the resulting data of the scan phase of a Chargeback Job.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	NUMERIC(19,0)	The internal ID of the Job run producing this data. This Id maps to the field ID of the CBJobDefinition table.
<b>Server</b>	NVARCHAR(16)	The NSS server that performed the scan for this data. Note that there could be several scanning servers in case of a multi-server Job, but there will be only one server performing the calculation [the master server].
<b>Location</b>	NVARCHAR(261)	Describes where the scan took place. It could be a path, for a file system scan, or a server name for an MS-Exchange or Quota scan.
<b>BillingClass</b>	INT	Depending on the Location and using the data from the CBM-PathMap table, this field will hold the Id of the BillingClass that applies to the Location.
<b>BillingAccount</b>	NVARCHAR(255)	The name of the Account to bill. This value depend of the AccountType filed of the CBJobDefintion table as follow: - AccountType = 0, BillingAccount = Windows User Account - AccountType = 1, BillingAccount = Windows Primary Group - AccountType = 2, BillingAccount = Path - AccountType = 3, BillingAccount = Child Path

COLUMN	TYPE	DESCRIPTION
<b>BillingGroup</b>	NVARCHAR(255)	The name of the Group the BillingAccount belongs to. This value depends on the AccountType filed of the CBJobDefinition table as follow: - AccountType = 0, BillingGroup = Windows Primary Group - AccountType = 1, BillingGroup = Domain - AccountType = 2, BillingGroup = Server's path - AccountType = 3, BillingGroup = Root path
<b>Value</b>	NUMERIC(19,0)	The actual usage retrieved during the scanning for this BillingAccount at this Location.
<b>Cost</b>	MONEY	The cost for the usage according to the Billing Class on the current Location. This is an intermediate calculation and the final cost is stored in the CBJobDataResult table. See below for more details on the cost calculations.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBJobData (ID numeric(19,0), Server nvarchar(16)
Location nvarchar(261), BillingClass int, BillingAccount
nvarchar(255), BillingGroup nvarchar(255), Value numeric(19,0),
Cost money);
```

The CBJobData table includes both usage and cost data, but the cost data is not final. It is the cost of the local usage and not the accumulated cost of the Cost Center across the entire Job. For example, a Cost Center might have usage in two Locations with different Billing Classes. The CBJobData table will store the local cost at each Location, whereas the CBJobDataResult table will store the final cost [which might be different than the sum of the costs if both Locations use the same Billing Class].

#### Table CBJobDataResult

This table holds the resulting data of the calculation phase of a Chargeback Job. It's actually the data of CBJobData table but grouped by BillingClass/BillingAccount pairs.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	NUMERIC(19,0)	The internal ID of the Job run producing this data. This Id maps to the field ID of the CBJobDefinition table.
<b>BillingClass</b>	INT	One of the BillingClass of the Job result. There will be one row in the CBJobDataResult per BillingClass/BillingAccount pair.
<b>BillingAccount</b>	NVARCHAR(255)	One of the BillingAccount of the Job result. There will be one row in the CBJobDataResult per BillingClass/BillingAccount pair.
<b>BillingGroup</b>	NVARCHAR(255)	The name of the Group for the current BillingAccount.

COLUMN	TYPE	DESCRIPTION
<b>Value</b>	NUMERIC(19,0)	The actual usage for the BillingAccount summed over all Locations of the CBJobData table.
<b>Cost</b>	MONEY	The cost for the usage according to the Billing Class/Billing Account across all Locations of the Job. This is the final cost for the Billing Class/Billing Account pair.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBJobDataResult (ID numeric(19,0), BillingClass int,
BillingAccount nvarchar(255), BillingGroup nvarchar(255), Value
numeric(19,0), Cost money);
```

The CBJobDataResult table includes both usage and cost data. The usage data is merely the sum of the usage at all Locations for a given Billing Class/Billing Account pair. The cost is the result of the Billing Class formula applied to the sum of usage of the Billing Account at all Locations using that Billing Class.

#### Table CBReportTopXDEF

This table holds the definition for all "Fastest Growing Cost Centers" Projections, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row per Projection run.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	INT	The internal ID of the Projection run.
<b>DataSrc</b>	TINY INT	The type of the data to used by the source Chargeback Jobs: - 0 = File System source - 1 = Northern Quota Server's quotas - 2 = MS Exchange mailboxes
<b>AccountType</b>	TINYINT	The type of the Cost Center of the source Jobs: - 0 = Windows User Account - 1 = Windows Primary User Group - 2 = Path/Location - 3 = Autopath, i.e. all first-level subfolders of a given path [only apply to File System data sources].
<b>GroupBy</b>	TINYINT	Specify how the data will be accumulated during the calculation : - 0 = by Billing Account [i.e. no accumulation] - 1 = by Billing Group - 2 = by Billing Class

COLUMN	TYPE	DESCRIPTION
<b>ValueType</b>	TINYINT	The type of the quantity used as usage in the source Chargeback Jobs: - 0 = usage is in size of data [e.g. size of the files] - 1 = usage is in count of data [e.g. count of files] The value 1 is reserved for future use.
<b>AnalyzeFrom</b>	DATETIME	The oldest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>AnalyzeTo</b>	DATETIME	The newest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>Runtime</b>	DATETIME	The date and time the Projection was run.
<b>Limit</b>	INT	The maximum number of Cost Centers to return. This value is entered by the user when the Projection is defined [e.g. Top 50 Fastest Growing Cost Centers].
<b>Definition</b>	VARBINARY (8000)	The internal Projection definition used by the NSS Core Service in binary form. This field is reserved for future use.
<b>UserID</b>	NUMERIC(19,0)	An internal ID used by NSS Core Service for referencing the Projection definition.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportTopXDEF (ID INT, DataSrc TINYINT, AccountType TINYINT, GroupBy TINYINT, ValueType TINYINT, AnalyzeFrom DATETIME, AnalyzeTo DATETIME, Runtime DATETIME, Limit int, Definition NVARCHAR(255), UserID numeric(19,0));
```

A Projection is calculated from the data of previously run Chargeback Jobs, no additional scan of the data source is made to complete the Projection. To determine which Jobs to include in the Projection, the Storage Chargeback engine uses the fields DataSrc, AccountType and ValueType as keys to the CBJobData table. This will return the IDs of all Chargeback Jobs that match the Projection definition [so that one does not get a Projection for File System with data from MS Exchange Jobs]. The fields AnalyzeFrom and AnalyzeTo are used filter out the data to specific date range.

#### **Table CBReportTopXRST**

This table holds the resulting data for all "Fastest Growing Cost Centers" Projection runs, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row

for each unique Group of each Projection run. A Group is defined in the Projection definition as either a single Billing Account, a Billing Group or a Billing Class.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	INT	The internal ID of the Projection run. This Id maps to the field ID of the CBReportTopDEF table.
<b>Field1</b>	NVARCHAR(255)	The name of one unique Group, e.g. a user account, a Billing Class name, etc...
<b>Field2</b>	FLOAT	The calculated growth for that Group is that Projection run. A value of 1 means no growth (100% of previous usage). A value of 2 means double growth (200% of previous usage).

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportTopXRST(ID INT, Field1 NVARCHAR(255), Field2
  FLOAT) ;
```

### Table CBReportProjCostDEF

This table holds the definition for all "Cost Prognosis per Cost Center" Projections, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row per Projection run.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	INT	The internal ID of the Projection run.
<b>DataSrc</b>	TINY INT	The type of the data to be used by the source Chargeback Jobs: - 0 = File System source - 1 = Northern Quota Server's quotas - 2 = MS Exchange mailboxes
<b>AccountType</b>	TINYINT	The type of the Cost Center of the source Jobs: - 0 = Windows User Account - 1 = Windows Primary User Group - 2 = Path/Location - 3 = Autopath, i.e. all first-level subfolders of a given path [only apply to File System data sources].
<b>GroupBy</b>	TINYINT	Specifies how the data will be accumulated during the calculation : - 0 = by Billing Account [i.e. no accumulation] - 1 = by Billing Group - 2 = by Billing Class
<b>ValueType</b>	TINYINT	The type of the quantity used as usage in the source Chargeback Jobs: - 0 = usage is in size of data [e.g. size of the files] - 1 = usage is in count of data [e.g. count of files] The value 1 is reserved for future use.

COLUMN	TYPE	DESCRIPTION
<b>AnalyzeFrom</b>	DATETIME	The oldest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>AnalyzeTo</b>	DATETIME	The newest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>Runtime</b>	DATETIME	The date and time the Projection was run.
<b>Date</b>	DATETIME	The date for which to calculate the projected cost. The Projection determines how much each Cost Centers will cost at the given date. This value is user-defined and is entered when the Projection is created.
<b>Definition</b>	VARBINARY (8000)	The internal Projection definition used by the NSS Core Service in binary form. This field is reserved for future use.
<b>UserID</b>	NUMERIC(19,0)	An internal ID used by NSS Core Service for referencing the Projection definition.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportProjCostDEF (ID INT, DataSrc TINYINT,
  AccountType TINYINT, GroupBy TINYINT, ValueType TINYINT,
  AnalyzeFrom DATETIME, AnalyzeTo DATETIME, Runtime DATETIME,
  Date DATETIME, Definition NVARCHAR(255), UserID numeric(19,0));
```

A Projection is calculated from the data of previously run Chargeback Jobs, no additional scan of the data source is made to complete the Projection. To determine which Jobs to include in the Projection, the Storage Chargeback engine uses the fields DataSrc, AccountType and ValueType as keys to the CBJobData table. This will return the IDs of all Chargeback Jobs that match the Projection definition [so that one does not get a Projection for File System with data from MS Exchange Jobs]. The fields AnalyzeFrom and AnalyzeTo are used filter out the data to specific date range.

#### Table CBReportProjCostRST

This table holds the resulting data for all "Cost Prognosis per Cost Center" Projection runs, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row for each unique Group of each Projection run. A Group is defined in the Projection definition as either a single Billing Account, a Billing Group or a Billing Class.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	INT	The internal ID of the Projection run. This Id maps to the field ID of the CBReportTopDEF table.
<b>Field1</b>	NVARCHAR(255)	The name of one unique Group, e.g. a user account, a Billing Class name, etc...

COLUMN	TYPE	DESCRIPTION
Field2	MONEY	The calculated cost projected at the specified date [field Date of the definition table].

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportProjCostRST(ID INT, Field1 NVARCHAR(255),
  Field2 MONEY);
```

#### Table CBReportProjCostAccDEF

This table holds the definition for all "Cost Prognosis per Cost Center Accumulated" Projections, regardless of the data source [file system, MS Exchange or Quotas]. This table is, in all aspects, similar to the CBReportProjCostDEF. The table definition and usage is exactly the same.

#### Table CBReportProjCostAccRST

This table holds the resulting data for all "Cost Prognosis per Cost Center Accumulated" Projections, regardless of the data source [file system, MS Exchange or Quotas]. This table is, in all aspects, similar to the CBReportProjCostDEF. The table definition and usage is exactly the same, except that Field2 will not include a singular cost at the projected date, but the accumulated cost at that date.

The cost is accumulated for each "Pay Day". The Pay Day is a time span [which can be more than one day] defined when the Projection is created, usually a week, a month or a year. The Projection engine extrapolates the usage of each Cost Center for each Pay Day, then calculates the cost of that usage according to the Billing Classes. Finally, it sums all calculates costs to return the final accumulated cost per Cost Center.

The Pay Day information is not publicly stored in the database. It is used only by the Projection engine when running the Projection. It is stored internally by the NSS Core Service inside the local XML file as well as in the binary Definition field.

#### Table CBReportProjDateDEF

This table holds the definition for all "End of Budget Prognosis per Cost Center" Projections, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row per Projection run.

COLUMN	TYPE	DESCRIPTION
ID	INT	The internal ID of the Projection run.
DataSrc	TINY INT	The type of the data to be used by the source Chargeback Jobs: - 0 = File System source - 1 = Northern Quota Server's quotas - 2 = MS Exchange mailboxes

COLUMN	TYPE	DESCRIPTION
<b>AccountType</b>	TINYINT	The type of the Cost Center of the source Jobs: - 0 = Windows User Account - 1 = Windows Primary User Group - 2 = Path/Location - 3 = Autopath, i.e. all first-level subfolders of a given path [only apply to File System data sources].
<b>GroupBy</b>	TINYINT	Specifies how the data will be accumulated during the calculation : - 0 = by Billing Account [i.e. no accumulation] - 1 = by Billing Group - 2 = by Billing Class
<b>ValueType</b>	TINYINT	The type of the quantity used as usage in the source Chargeback Jobs: - 0 = usage is in size of data [e.g. size of the files] - 1 = usage is in count of data [e.g. count of files] The value 1 is reserved for future use.
<b>AnalyzeFrom</b>	DATETIME	The oldest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>AnalyzeTo</b>	DATETIME	The newest running date for a Chargeback Job to be included in this Projection. Only Jobs with running dates between AnalyzeFrom and AnalyzeTo value are used.
<b>Runtime</b>	DATETIME	The date and time the Projection was run.
<b>Cost</b>	MONEY	The target cost for which to calculate the projected date. The Projection determines at which date each Cost Centers will reach this cost. This value is user-defined and is entered when the Projection is created.
<b>Definition</b>	VARBINARY (8000)	The internal Projection definition used by the NSS Core Service in binary form. This field is reserved for future use.
<b>UserID</b>	NUMERIC(19,0)	An internal ID used by NSS Core Service for referencing the Projection definition.

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportProjDateDEF (ID INT, DataSrc TINYINT,
AccountType TINYINT, GroupBy TINYINT, ValueType TINYINT,
AnalyzeFrom DATETIME, AnalyzeTo DATETIME, Runtime DATETIME,
Cost MONEY, Definition NVARCHAR(255), UserID numeric(19,0));
```

A Projection is calculated from the data of previously run Chargeback Jobs, no additional scan of the data source is made to complete the Projection. To determine which Jobs to include in the Projection, the Storage Chargeback engine uses the fields DataSrc, AccountType and ValueType as keys to the CBJobData table. This will return the IDs of all Chargeback Jobs that

match the Projection definition [so that one does not get a Projection for File System with data from MS Exchange Jobs]. The fields AnalyseFrom and AnalyseTo are used filter out the data to specific date range.

**Table CBReportProjCostRST**

This table holds the resulting data for all "End of Budget Prognosis per Cost Center" Projection runs, regardless of the data source [file system, MS Exchange or Quotas]. The table will hold one row for each unique Group of each Projection run. A Group is defined in the Projection definition as either a single Billing Account, a Billing Group or a Billing Class.

COLUMN	TYPE	DESCRIPTION
<b>ID</b>	INT	The internal ID of the Projection run. This Id maps to the field ID of the CBReportTopDEF table.
<b>Field1</b>	NVARCHAR(255)	The name of one unique Group, e.g. a user account, a Billing Class name, etc...
<b>Field2</b>	DATE	The calculated date when the Cost Center will reach the given cost [field Cost of the definition table].

The NSS Core Service creates this table using the SQL statement:

```
CREATE TABLE CBReportProjDateRST(ID INT, Field1 NVARCHAR(255),
Field2 DATE);
```

**Table CBReportProjDateAccDEF**

This table holds the definition for all "End of Budget Prognosis per Cost Center Accumulated" Projections, regardless of the data source [file system, MS Exchange or Quotas]. This table is, in all aspects, similar to the CBReportProjDateDEF. The table definition and usage is exactly the same.

**Table CBReportProjDateAccRST**

This table holds the resulting data for all "End of Budget Prognosis per Cost Center Accumulated" Projections, regardless of the data source [file system, MS Exchange or Quotas]. This table is, in all aspect, similar to the CBReportProjDateDEF. The table definition and usage is exactly the same, except that the date returned in Field2 will be the date when the Cost Center's accumulated cost will reach the given target.

A Cost Center's cost is accumulated for each "Pay Day". The Pay Day is a time span [which can be more than one day] defined when the Projection is created, usually a week, a month or a year. The Projection engine extrapolates the usage of each Cost Center for each Pay Day, then calculates the cost of that usage according to the Billing Classes. It then sums those costs and returns the date when this accumulated cost equals or exceeds the given target cost.

The Pay Day information is not publicly stored in the database. It used only by the Projection engine when running the Projection. It's stored internally by the NSS Core Service inside the local XML files as well as in the binary Definition field.


## ABOUT PROJECTION EXTRAPOLATION

All Projections use the same formula for extrapolating cost trends. The formula is used for both exponential or linear extrapolation.

Only the Fastest Growing Cost Center uses the linear extrapolation since the growth speed used in the calculation can be both positive and negative [negative numbers cannot be used in the exponential extrapolation].


All other Projections use the exponential extrapolation. This gives much more accurate results than the typical linear extrapolation. The difference in accuracy between the linear and exponential models are shown in the two graphs below:

Figure 71: Error margin between linear and exponential extrapolation for an accelerating trend



For an accelerating trend, the linear model shows a much lower increase than what could be expected.

Figure 72: Error margin between linear and exponential extrapolation for slowing trend.



For an slowing trend, the linear model shows a much higher increase than what could be expected.

The actual formula is as follows:

$$Y = \frac{\sum_{i=1}^n x_i^2 \cdot \sum_{i=1}^n y_i - \sum_{i=1}^n x_i \cdot \sum_{i=1}^n x_i y_i + \left( n \cdot \sum_{i=1}^n x_i y_i - \sum_{i=1}^n x_i \cdot \sum_{i=1}^n y_i \right) \cdot X}{n \cdot \sum_{i=1}^n x_i^2 - \left( \sum_{i=1}^n x_i \right)^2}$$

where:

- X is the target date of the Projection
- For Fastest Growing, Y is the projected growth,
- For other Projections, Y is the natural logarithm of the projected size. So the projected usage is:

$$ProjectedUsage = e^Y$$

- $x_1, x_2, \dots, x_n$  are dates when data is available,
- For Fastest Growing,  $y_1, y_2, \dots, y_n$  are the growth speed at the date  $x_i$
- For other Projections,  $y_1, y_2, \dots, y_n$  are natural logarithm of storage usage at the date  $x_i$ ,

$$y_i = \ln(usage(x_i))$$



## ABOUT MSDE

Northern Storage Suite is available with Microsoft Database Engine [MSDE]. MSDE is a limited version of Microsoft SQL. It is limited to five database connections and its databases are limited to 2 GB in size. Each instance of Northern Storage Suite client will use one such connection.

It should be kept in mind that MSDE is by no means an ideal database. It has the aforementioned limitations and furthermore is limited with respect to through-put and overall performance. Northern provides this option to ease evaluation only.

Severe conflicts will occur if MSDE is installed on a computer running MS SQL Server. A warning will be provided by the Northern Storage Suite installer if such an installation is attempted.

## CONFIGURING MS SQL SERVER

Two configuration steps are required to be able to use MS-SQL Server.

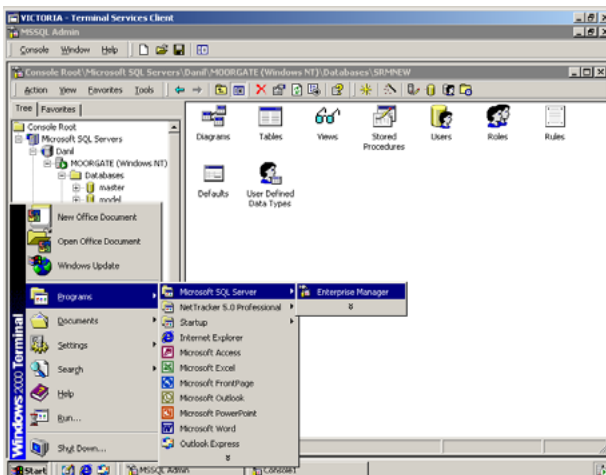
- Creating a new database
- Setting proper database permissions for the NSS Service

## CREATING A NEW DATABASE IN MS SQL

To create a new MS SQL database:

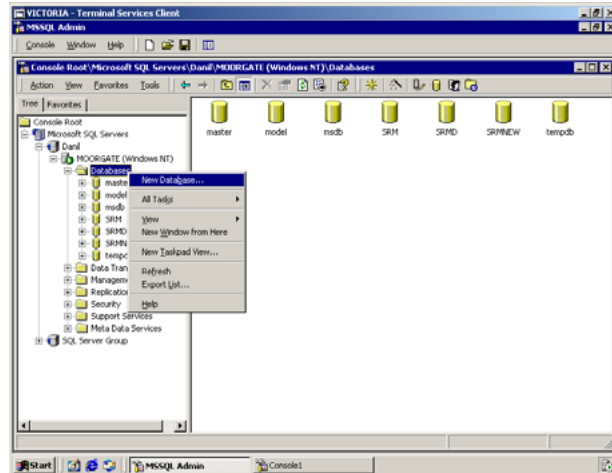
- 1 Choose: Start > Programs > Microsoft SQL Server > Enterprise Manager.

Figure 73: Starting Microsoft SQL



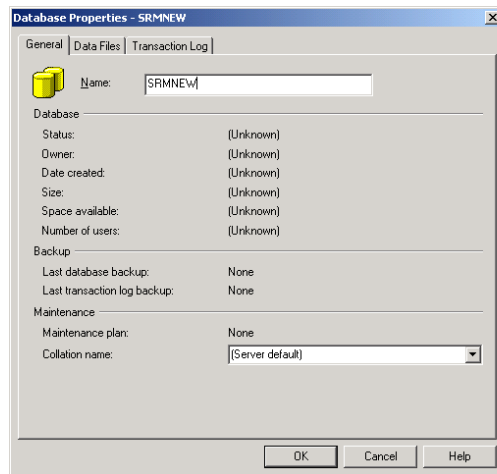
- 2 In the MS SQL Manager, expand your desired SQL server and right-click the Databases directory and select “New database...” from the shortcut menu. The Database Properties dialog appears.

Figure 74: Creating a new database in MS SQL



- 3 Give the new database a name, and click OK.

Figure 75: Database Properties dialog box



- 4 Give the account permissions to the database

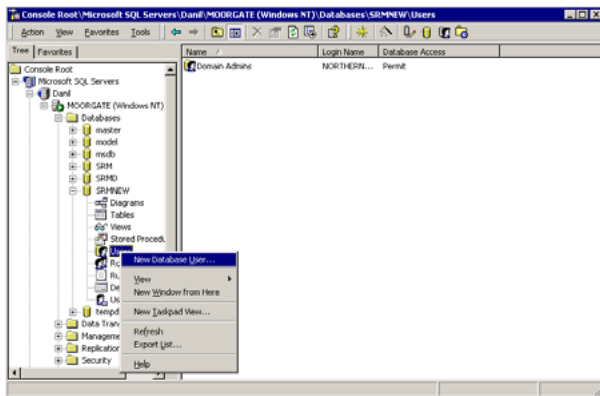
## SETTING DATABASE PERMISSIONS IN MS SQL

The data stored in the database will be both read and written by the Northern Storage Suite service. Those operations will only work if the account running the NSS Service has full access to the database.

To create an SQL account for the Northern Storage Suite Service:

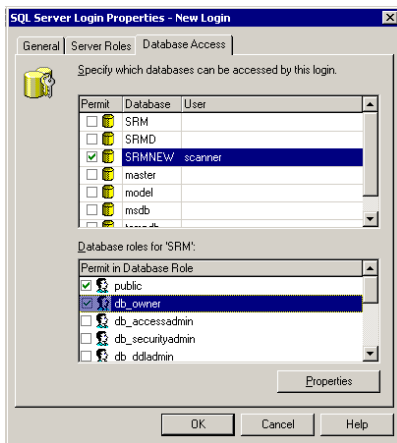
- 1 Right-click the Users object in the database you have just created, and choose “New database User...” from the shortcut menu. The New Login dialog appears.

Figure 76: Creating a new database Use



- 2 In the New Login dialog, on the Database Access tab, make sure that the Northern Storage Suite Service's account is given the database role “db\_owner”.

Figure 77: SQL Server Login Properties - New Login dialog box appears



- 3 Click OK



# NORTHERN

## NORTHERN: THE COMPANY

Northern Parklife, Inc. [NORTHERN] develops storage management software for Microsoft Windows server environments. Working with a global network of resellers and authorized distributors, the company's products are sold in 52 countries worldwide. NORTHERN® supplies a global market through its two offices, Tampa [Florida, US] and Stockholm [Sweden].

Northern is distinguished by the flexibility and reliability of its family of storage management solutions, as well as its unmatched customer support and responsiveness to market demand.

### BUSINESS CONCEPT

NORTHERN develops and simplifies storage administration in Microsoft Windows server environments with advanced software technologies. Our products and employees help systems administrators to succeed in rationalizing and gaining better control over their IT resources.

### Business Mission

NORTHERN's business mission is to provide system administrators with more **Time to Manage Storage**. The design core of all NORTHERN solutions is the capacity to re-balance the equation of increased responsibilities and static or diminishing resources. System administrators are saddled with an increasing range of responsibilities while their resources, particularly their time, remain limited. NORTHERN delivers a range of inherently tactical, while at the same time easy-to-use storage management tools that serve to augment the IT organization by reducing workload and increasing efficiencies. Our software makes storage administration simpler and less prone to error, and gives administrators greater latitude in developing and implementing creative business solutions that enhance their value to the organization.

NORTHERN software offers system administrators:

- More time for pro-active planning and creative problem-solving
- Better oversight and control
- Easy, hassle-free installations
- Simple, friction-free administration
- Reliability
- Fewer end-user issues (lightening the support burden)
- More spare time and free weekends

**Time to Manage Storage** is the removal of unnecessary repetition and complicated word-arounds, providing time for pro-active Storage Resource Management and time for job fulfillment.

### THE MARKET

NORTHERN is the leader in the European systems administration software solutions market, which is comprised of systems administrators and network managers tasked with managing a broad range of network responsibilities. Because networking issues know no geographical

boundaries, the European and US markets served by NORTHERN allow for a consistent global marketing and product development strategy.

## **PRODUCT DEVELOPMENT**

NORTHERN's product development is managed internally and driven by feedback from customers, Windows NT/2000 users and discussions with NORTHERN's board of advisors that consists of financial, technical and market experts.

## **CUSTOMER SUPPORT PROGRAM**

Take advantage of expert technical guidance and inexpensive upgrades with NORTHERN's Customer Support Program

You have invested in a CD-ROM from NORTHERN, a software product. Every effort has been made to ensure that this software is your solution. But sometimes something extra is needed, to guarantee your investment.

The need for continuous organizational advances, network configurations, hardware peculiarities; these factors can lead to time consuming issues and the need for constant system evolution. This is when the NORTHERN Customer Support Program becomes invaluable: uncomplicated and inexpensive upgrading of your system to the most advanced solutions available, and smooth installation and integration with expert technical guidance.

The NORTHERN Customer Support Program brings you closer to the technological front line and provides you with the assistance you need to make the most out of being there.

NORTHERN does not just make software - we provide our customers with complete solutions.

## **FEATURES**

NORTHERN's Customer Support Program offers two different kinds of support: Premium and Standard.

More information about NORTHERN's Customer Support Program is available at [www.northern.net/support](http://www.northern.net/support).

Please contact your local software supplier to find out how you can join NORTHERN's Customer Support Program. A list of authorized NORTHERN resellers is available at [www.northern.net](http://www.northern.net).

## **CONTACTING NORTHERN**

Do you have questions? Would you like to comment on our products, support or business? Do you have suggestions for future releases or products? Would you like to pay us a visit?

### **STREET ADDRESS**

Northern Parklife Inc  
5201 West Kennedy Blvd Suite#930  
Tampa FL 33609  
USA

Northern Parklife AB  
Gavlegatan 18b  
SE-113 30 Stockholm  
Sweden

### **PHONE**

US time zones: 1.800.881.4950  
1.813.639.0767

Other time zones: +46 8 457 50 00

### **FAX**

US time zones: 1.813.639.0757

Other time zones: +46 8 457 50 50

### **E-MAIL**

Sales inquiries: [sales@northern.net](mailto:sales@northern.net)  
Support questions: [support@northern.net](mailto:support@northern.net)  
Other: [info@northern.net](mailto:info@northern.net)

### **WEB**

[www.northern.net](http://www.northern.net)





