Configuring

Contents

Configuring

The topics in this section provide information related to configuring your DS8000®. Topics include storage and network configuration, host attachment, and logical partition.

1. Creating a custom logical storage configuration

Complete this task to create a custom logical configuration where you manually determine the storage allocation details. You must complete all of the steps and substeps in this task to create your configuration.

Before you begin

You must have installed and configured the DS Storage Manager before you can create your custom logical configuration. You also must have obtained and applied an Operating Environment License (OEL) with a maximum storage amount that exceeds the amount of storage that you intend to configure. In addition, your network must be configured so that it can support all the components and functions that you will use with your storage unit.

About this task

You can create a logical configuration by following the steps in this task.

Note: Before you configure, disable any firewalls, as they might interfere with DS8000® communication.

Return to the next step of this task after completing each of the subtasks below.

Procedure

1. Create your arrays. To create an array, you must complete the subtask in Creating arrays.
2. Create your ranks. To create a rank, you must complete the subtask in Creating ranks.
3. Create your extent pools. To create an extent pool, you must complete the subtask in Creating extent pools (automatic mode) or Creating extent pools (manual mode).
4. Perform one of the following tasks for either your fixed block (FB) or CKD host:
   - If you are working with a CKD host, you must create CKD LCUs and volumes:
     - To create LCUs and volumes, complete the steps in Creating CKD LCUs with volumes.
   - If you are working with an FB host, you must create FB volumes and volume groups:
     - To create volumes, complete the steps in Creating FB volumes.
     - To create volume groups, complete the steps in Creating FB volume groups.
5. Reconfigure your host system to start sending I/O to the DS8000. To reconfigure your host system, complete the steps in Enabling host system I/O.

Creating arrays
Complete this task to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

Creating ranks
Complete this task to create ranks and to specify their arrays and extent pools.

Creating extent pools (automatic mode)
Complete this task to create an extent pool using the automatic configuration mode that allows you to rely on automatic system calculations for much of the configuration.

Creating extent pools (manual mode)
Complete this task to create an extent pool using the manual configuration mode.

Creating CKD LCUs with volumes
Complete this task to create CKD logical control units (LCUs) and volumes and to specify their attributes and properties.

Creating FB volumes
Complete this task to create FB (fixed block) volumes for open systems and to specify their attributes and properties.

Creating FB volume groups

Complete this task to create FB (fixed block) volume groups for open systems and to specify their properties and parameters.

Enabling host system I/O
Complete this task to perform the high-level steps that are required to enable your host to send I/O to the DS8000.

Related concepts
Getting started with configurations
Activating licensed functions
Logical configuration overview

Related tasks
Configuring I/O ports
Creating new hosts
Creating arrays
Creating CKD LCUs with volumes
Creating FB volumes
Creating FB volume groups

Related reference
Create new ranks

Related information
Running the DS CLI

1.1. Creating arrays

Complete this task to create arrays, either automatic or custom, and to specify their RAID and rank attributes.

About this task

Perform the following steps to create arrays:

Procedure

1. In the navigation, select **Pools** > **Internal Storage**. On the Internal Storage page underneath Manage Internal Storage, click the Arrays tab.
2. From the **Action** drop-down menu, select **Create Arrays**. The Create Array Definition method page is displayed.
3. On the Create Array – Definition method page, choose to either create your arrays automatically, where the DS Storage Manager specifies the array sites, or to create custom arrays where you specify the array sites. Click **Next**.
   - If you chose to create your arrays automatically, in Create Array – Array configuration (Auto), specify the quantity and RAID type of the arrays that you are creating. Click **Next** after you have made your selections. The Create Array – Add arrays to ranks page is displayed.
   - If you chose to create custom arrays, in Create Array – Array configuration (Custom), specify the RAID type for the arrays that you are creating and select the array site numbers that you want to assign to the arrays. Click **Next** after you have made your selections. The Create Array – Add arrays to ranks page is displayed.
4. In Create Array – Add arrays to ranks, you can optionally choose to add the arrays that you are creating to ranks. If you choose to add your arrays to ranks, you must specify the storage type, FB or CKD, for the ranks that you are creating. Click **Next**. The Create Array – Verification page is displayed.
5. In Create Array – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the array creation process.

What to do next

If you are creating a new logical storage configuration, return to Creating a custom logical storage configuration and complete the rest of the steps.

Parent topic: Creating a custom logical storage configuration

Related reference
Create array — Add arrays to rank
Single array properties — Status

1.2. Creating ranks

Complete this task to create ranks and to specify their arrays and extent pools.

About this task

Perform the following steps to create ranks:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. On the Internal Storage page underneath Manage Internal Storage, click the Ranks tab.
2. From the **Action** list, select **Create...**. The Create Rank – Select array for rank page is displayed.
3. Select an array from the table. Then click **Next**. The Create Rank – Define rank properties page is displayed.

   Note: If there are no available arrays for selection, you can create one by clicking the **Create new array** button. When you have finished creating an array, you are returned to this page, and the array that you created is available for selection.
4. Specify the **Storage type** and then click **Next**. The Create Rank – Select extent pool page is displayed.
5. Optionally, you can select one extent pool. Then click **Next**. The Create Rank – Verification page is displayed.

   Note: If there are no available extent pools for selection, you can create one by clicking the **Create new extent pool** button. When you have finished creating an extent pool, you are returned to this page, and the extent pool that you created is available for selection.
6. In Create Rank – Verification, review the attributes and values to verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and then specify the correct values. Otherwise, click **Finish** to complete the rank creation process.

What to do next

If you are creating a new logical storage configuration, return to Creating a custom logical storage configuration and complete the rest of the steps.

Parent topic: Creating a custom logical storage configuration

1.3. Creating extent pools (automatic mode)

Complete this task to create an extent pool using the automatic configuration mode that allows you to rely on automatic system calculations for much of the configuration.

About this task

The automatic configuration mode allows the system to calculate and assign values after you have supplied key data. The results of combining your input and the system calculation are displayed within a bar graph on the Create Extent Pools page. You can use the bar graph to determine whether the extent pool configuration meets your needs.

Perform the following steps to create extent pools:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. On the Internal Storage main page, select the Extent Pools tab. Select **Create extent pools . . .** from the drop-down list.

   The Create Extent Pools page is displayed with **Automatic** shown as the Type of Configuration. Do not change this setting.

2. Supply the following information:

   Storage Type

   Select one of the following storage types:

   - FB: Use the FB storage type for open systems host volumes. Specifies that the storage type is fixed block extents - 1 GB. In fixed block architecture, the data (the logical volumes) are mapped over fixed-size blocks or sectors.
   - CKD: Use the CKD storage type for System z® host volumes. Specifies that the storage type is count-key-data extents - CKD Mod 1. In count-key-data architecture, the data field stores the user data.

   RAID Type

   Select one of the following RAID types:

   - RAID 5
   - RAID 6
   - RAID 10

   RAID type is used by the ranks in the extent pool. RAID is a methodology for grouping disk drives to protect against data loss from a failing disk drive.

   Type of Configuration

   Leave the selection set to Automatic.

   DA Pair Usage

   Select one of the following from the pull-down list:

   - Spread among all pairs
   - Spread among least used pairs
   - Sequentially fill all pairs

   The bar graph displays in real-time the results of your choice. This information is used to analyze the best configuration selection for your system.

   Encryption Group

   If encryption is used, select the encryption group ID that will be associated with the rank. If encryption is not used, `None` is the only option. This field is enabled only if no ranks have been created or no ranks are in the process of being created.

   Drive Class

Select one of the disk drive modules in the pull-down list. If one of the selections within this field is *146 GB 15K Enterprise Drive with xxxx GB available*, this example represents the following designation:

- The disk drive module is a high-speed Fibre Channel disk drive.
- The disk drive module has a size of 146 GB and a speed of 15,000 RPM.

When you make your selection, an information box is displayed that lists the available capacity as ranks and the capacity size that is needed to create the extent pool.

Select Capacity to Configure

Select the appropriate rank designation for the capacity that you want to configure. The bar graph displays the results of the choice that you have selected.

Note: As soon as you supply these values, the results of your selections are immediately reflected in the bar graph and the Select Capacity to Configure list is updated.

3. Make one of the following selections for Number of Extent Pools within the Define Extent Pool Characteristics section of the page:

   Two extent pools (one per server)

   Causes the system to assign the capacity to two extent pools, which makes a pair.

   Extent pool for each rank (physical isolation)

   Causes the system to assign one array to each extent pool and evenly assign the extent pools to each server if the Automatic option is in effect under Server Assignment.

   Single extent pool

   Causes the system to assign the capacity to a single extent pool.

4. Enter an extent pool name prefix consisting of 1 to 12 characters as a value for the Nickname prefix field.

   The prefix allows you to supply a unique identifier for the extent pool.

5. Select the type of suffix from the Suffix drop-down box.

   - None
   - ID
   - Custom

   If you select `ID`, the system automatically adds the next available numeric suffix to the prefix starting with "_0".

6. Either leave the default values for Storage Reserved and Storage Threshold or enter new values 0 - 100. These percentage values are indicators that the system uses to issue an alert when the value is exceeded.

   Notes:

   a. The storage reserved value represents the percentage of extents in the extent pool that are reserved and cannot be allocated for use.

   b. The storage threshold value represents a percentage of allocated extents out of the total nonreserved storage extents. The system generates an alert when the threshold is reached. It allows you time to make adjustments before a storage full condition occurs.

7. Select the server assignment from the Server Assignment drop-down list.

   You can leave the selection on Automatic and the system assigns the extent pool to the server (0 or 1). That is the most efficient way to maintain server balance. However, you can control the server assignment by selecting server 0 or server 1.

   Note: The server assignment remains Automatic and cannot be changed if the Two Extent Pools (1 per server) option is selected under Number of Extent Pools.

As soon as you make your server choice, the server information area is updated. This information can help you balance the storage server workload.

8. Click **OK** if you are ready to create the extent pools you have selected. You are returned to the updated Create Extent Pool Verification table. Click **Add Another Pool** if you are creating more extent pools and the Create Extent Pool Verification table is updated while the form remains available to create another extent pool.
9. Review the changes to the Create Extent Pool Verification table. If you are satisfied with the changes, click **Create All** to accept the changes or click **Cancel** to cancel out of all of the changes.

   While still at the Create Extent Pool Verification table, the following actions are available from the table drop-down list:

   Add Pools

   You can create additional new extent pools from the Create New Extent Pools page.

   Add Capacity to Pool

   You can add capacity to a pool that is selected in the Create Extent Pool Verification table.

   Delete Pools

   You can delete one or more pools that are selected in the Create Extent Pool Verification table.

Parent topic: Creating a custom logical storage configuration

Related concepts
RAID 5 overview
RAID 10 overview

Related tasks
Creating extent pools (manual mode)

Related reference
Create extent pool verification table
Create extent pools

1.4. Creating extent pools (manual mode)

Complete this task to create an extent pool using the manual configuration mode.

About this task

The manual configuration mode allows you to create an extent pool and have control over the array site, array, rank assignment, and the server assignment for the extent pool.

When you select **Manual** for the Configuration Type option on the Create Extent Pool page, the page does not change. Only the Select Available Capacity section of the page changes by displaying an Available components table that includes: array sites, arrays, and rank assignments. When you select a component that you want to use to create your extent pool, the bar graph updates in real-time. You can use the bar graph to determine whether the extent pool configuration meets your requirements.

Perform the following steps to create extent pools:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. On the Internal Storage main page, select the Extent Pools tab. Select **Create extent pools . . .** from the drop-down list.

   The Create Extent Pools page is displayed with **Automatic** shown as the Type of Configuration.
2. Select the applicable storage type and RAID type from the drop-down lists in the Define Storage Characteristics section of the page. Use the RAID pull-down list to select one of the following values:
   - RAID 5
   - RAID 6
   - RAID 10

   RAID type is used by the ranks in the extent pool. RAID is a methodology for grouping disk drives to protect against data loss from a failing disk drive. An extent pool can contain ranks of different RAID types. An extent pool of this type is listed with a RAID type of mixed on the Extent Pools tab of the Internal Storage main page.
3. Select **Manual** from the Type of Configuration drop-down list. The Select Available Capacity section of the Create Extent Pools page changes format and displays a table of available arrays, array sites, and ranks that can be used to create the extent pool.
4. Select one or more of the available components from the table. The results of your selection are immediately reflected in the bar graph. If there are no available components in the table, try changing the RAID Type and or the Storage Type.
5. Make one of the following selections for Number of Extent Pools within the Define Extent Pool Characteristics section of the page:

   Two extent pools (one per server)

   Causes the system to assign the capacity to two extent pools, which makes a pair.

   Extent pool for each rank (physical isolation)

   Causes the system to assign one array to each extent pool and evenly assign the extent pools to each server.

   Single extent pool

   Causes the system to assign the capacity to a single extent pool.
6. Enter an extent pool name prefix consisting of 1 to 12 characters as a value for the Nickname prefix field.

   The prefix allows you to supply a unique identifier for the extent pool.
7. Select the type of suffix from the Suffix drop-down box.
   - None
   - ID

- Custom

If you select `ID`, the system automatically adds the next available numeric suffix to the prefix starting with "_0".

8. Either leave the default values for Storage Reserved and Storage Threshold or enter new values 0 - 100. These percentage values are indicators that the system uses to issue an alert when the value is exceeded.

   Notes:

   a. The storage reserved value represents the percentage of extents in the extent pool that are reserved and cannot be allocated for use.

   b. The storage threshold value represents a percentage of allocated extents out of the total non-reserved storage extents. The system generates an alert when the threshold is reached. It allows you time to make adjustments before a storage full condition occurs.

9. Select the server assignment from the Server Assignment drop-down list.

   You can leave the selection on Automatic and the system assigns the extent pool to the server (0 or 1) that it determines is the most efficient to maintain server balance. However, you can control the server assignment by selecting server 0 or server 1.

   Note: The server assignment remains Automatic and cannot be changed if the Two Extent Pools (1 per server) option is selected under Number of Extent Pools.

   As soon as you make your server choice, the server information area is updated. This information can help you balance the storage server workload.

10. Click **OK** if you are ready to create the extent pools you have selected. You are returned to the updated Create Extent Pool Verification table. Click **Add Another Pool** if you are creating another extent pool and the Create Extent Pool Verification table is updated while the form remains available to create another extent pool.
11. Review the changes to the Create Extent Pool Verification table. If you are satisfied with the changes, click **Create All** to accept the changes or click **Cancel** to cancel out of all of the changes.

    While still at the Create Extent Pool Verification table, the following actions are available from the table drop-down list:

    Add Pools

    You can create additional new extent pools from the Create New Extent Pools page.

    Add Capacity to Pool

    You can add capacity to a pool that is selected in the Create Extent Pool Verification table.

    Delete Pools

    You can delete one or more pools that are selected in the Create Extent Pool Verification table.

Parent topic: Creating a custom logical storage configuration

Related concepts
RAID 5 overview
RAID 10 overview

Related tasks
Creating extent pools (automatic mode)

Related reference
Create extent pool verification table
Create extent pools

1.5. Creating CKD LCUs with volumes

Complete this task to create CKD logical control units (LCUs) and volumes and to specify their attributes and properties.

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. On the CKD LCUs and Volumes – Main page, from the **Storage image** drop-down list, select a storage image.
2. In the Tasks section, click the **Create new LCUs with volumes** link. The Create LCUs – Define LCUs page is displayed.
3. Select one or more LCUs from the list of those available.

   Note: If you are creating LCUs to associate with a specific extent pool, you must select LCU ID numbers that match, as even or odd, the server number that was specified for the extent pool when it was created. For example, an extent pool that is associated with Server 1 can be associated only with LCUs that have odd numbers.

4. Define the parameters for the selected LCUs.
   a. Specify the LCU type.
   b. You can change the default SSID by highlighting and typing over it.
   c. You can enable critical mode to control the behavior of the PPRC pairs that have a primary logical devise on this LCU in a consistency group. If you enable this option, write operations to the source remote mirror and copy volume are not allowed if data cannot be copied to the target volume of the volume pair.

      Note: This option is available for administrators only.

   d. Establish the timeout times in seconds.
   e. You can enable a consistency group.
5. Click **Next** to continue. The Create Volumes – Define Base Volumes page is displayed.
6. Select the base type. If you select 3390 Mod A or custom, you have the option of specifying the volume size. Enter the quantity of base volumes. You also must specify the base start address, and a sorting order of Ascending or Descending. Select the storage allocation method. If you select standard, you must also specify the extent allocation method and the performance group number.
7. If you select to assign alias volumes to the base volumes, you must also specify the following items:
   a. Specify the alias starting address and address sorting order.
   b. Select a method for assigning aliases to bases.
      - If you evenly assign alias volumes among bases, then you must specify the total quantity of alias volumes.
      - If you assign aliases using a ratio of aliases to base volumes, then you must specify the ratio.
8. In the Assign nicknames to volumes section, you may choose to generate a sequence of nicknames based on the specified prefix and suffix and if you want to use hexadecimal values. If you are using hexadecimal values, see Using hexadecimal values for conversion information. To create another volume, click **Add another**. If you are finished, click **OK**. The Create Volumes table is displayed. From this table, you can create, modify or delete volumes. To add the selected volumes to the LCU, click **Next**. The LCU to Extent Pool Assignment page is displayed.
9. Select the assignment method and indicate if the LCU may span multiple pools. Click **Finish**. The Create LCUs Verification page is displayed.
10. Use the Verification page to review the established attributes and verify that they are correct. You can also add, modify or delete LCUs. Click **Create All** to complete the LCU creation process. The system creates the LCUs in increasing LCU ID order, one LCU and its volumes at a time.

What to do next

If you are creating a new logical storage configuration, return to Creating a custom logical storage configuration and complete the rest of the steps.

Parent topic: Creating a custom logical storage configuration

1.6. Creating FB volumes

Complete this task to create FB (fixed block) volumes for open systems and to specify their attributes and properties.

About this task

Note: Before you configure, disable any firewalls, as they might interfere with DS8000® communication.

To create FB volumes, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **FB Volumes**. The FB Volumes - Main page is displayed. Select the storage image from the **Storage image** drop-down list.
2. In the Tasks section, click the **Create new volumes** link. The Create Volumes - Select extent pools page is displayed.
3. You must select an extent pool for the target volumes. To ensure a balanced configuration, select extent pools in pairs (one from each server). If you select multiple pools, you must also select a volume assignment option.After you select the extent pool, click **Next**. The Add Volumes – Define volume characteristics page is displayed.
4. Define the volume type for the new volumes. When you specify the volume size and quantity, the projected remaining capacity is updated. Specify the storage allocation method. If you select Extent Space Efficient (ESE) or Standard, then you must specify the extent allocation method. If you select Standard, then you must also assign a performance group ID.
5. You may optionally choose a nickname for the volume, based on the specified prefix and suffix. Select **Hexadecimal sequence** to use hexadecimal values. If you are using hexadecimal values, see Using hexadecimal values for conversion information.

   Tip: If you plan to create volume groups, you can use a unique and meaningful nickname that can help you easily find the volumes that you want to include in a volume group. For example, if you are creating multiple volumes of data for a specific department, you can make the nickname prefix an abbreviation of that department's name. Then you can use a predetermined range of numbers in the nickname suffix to identify each individual volume.
6. You can select any number of volume groups from the volume groups list to associate with the new volumes.
7. When you have finished defining the volume characteristics, click **Add Another** to add another volume, or **OK** to display the Create volumes page.
8. On the Create volumes page you can add, modify or delete volumes. Select the volumes that you want to assign to the previously selected extent pools and click **Next**. The Create Volumes - Select LSS page is displayed.
9. Select the LSS assignment method. If you select Automatic, the system will assign the volume addresses. If you select Manual, select one or more LSSs from the list to assign volume addresses. Click **Finish**. The Create volumes verification page is displayed.
10. Review the attributes and values to verify that they are correct. If they are not correct for a specific volume, select **Modify** from the **Action** drop-down list to return to the appropriate page and change the values. To add additional volumes, select **Add volumes**. To remove a volume from the list, select **Delete**. When all of the volumes listed are ready, click **Create All** to complete the volume creation process.

What to do next

If you are creating a new logical storage configuration, return to Creating a custom logical storage configuration and complete the rest of the steps.

Parent topic: Creating a custom logical storage configuration

Related concepts
Open systems

1.7. Creating FB volume groups

Complete this task to create FB (fixed block) volume groups for open systems and to specify their properties and parameters.

Before you begin

Volume groups provide a volume masking function. When you create a volume, you can assign it to one or more volume groups. Fixed block (FB) or open systems server attachments (hosts) are assigned to volume groups to control access to the volumes. There can be up to 8320 volume groups. Any host port (or host port group) can access only one volume group. Normally, a host is mapped to only one volume group; however, some hosts such as Sun Solaris allow mapping to multiple volume groups. A volume group can be accessed by multiple servers(hosts) if they have the same block size and address. If there is a requirement for volume sharing, for example in a clustering environment, then a volume can be a member of multiple volume groups. In both cases, data integrity is the responsibility of host software.

About this task

To create FB volume groups, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **Volume Groups**. The Volume groups- Main page is displayed. Select the storage image from the **Storage image** drop-down list.
2. From the **Action** list, select **Create**. The Create New Volume Group – Define volume group properties page is displayed.
3. Define the properties. In the **Volume Group Nickname** field, you can either keep or modify the default nickname. In **Host Type**, select the appropriate host type. Both the **Nickname** and **Host Type** fields are required. Select the volumes that you want to include in the volume group. After you define the properties, click **Next**.
4. Use the Create Volume Group – Verification page to review the attributes and verify that they are correct. If the attributes and values are not correct, click **Back** as appropriate to return and to specify the correct values. Otherwise, click **Finish** to complete the volume group creation process.

Parent topic: Creating a custom logical storage configuration

Related concepts
Open systems

1.8. Enabling host system I/O

Complete this task to perform the high-level steps that are required to enable your host to send I/O to the DS8000®.

Before you begin

You must complete your logical storage configuration before you attempt to send host I/O to the DS8000.

About this task

This task contains the high-level tasks that you must complete before you can successfully send host I/O to your DS8000. Instructions and details for each individual step can be found in the host system documentation or the *IBM® System Storage DS8000 Host Systems Attachment Guide*.

Procedure

1. Verify and update the supported fabric switch firmware.
2. Create your storage area network (SAN) zones and assign them to the appropriate zone sets.
3. Verify and update the host adapter firmware on the host system.
4. Install, verify, and update the host failover driver (SDD) on the host system.
5. If necessary, reconfigure or reboot the host system.
6. Verify that the host system can access the DS8000 logical volumes that are assigned to it.

What to do next

If you are creating a new logical storage configuration, return to Creating a custom logical storage configuration and complete the rest of the steps.

Parent topic: Creating a custom logical storage configuration
Library | Support | Terms of use | Feedback

2. Managing logical storage configurations

The topics in this category present information related to working with the various elements of an existing logical configuration, such as arrays, ranks, extent pools, and volumes. Information on the pages that are associated with these tasks can be found in the Reference section.

Viewing and modifying properties
Complete this task to access properties pages so that you can review or modify the properties.

Increasing FB volume capacity
Complete this task to increase the volume capacity of a fixed block (FB) host system.

Increasing CKD volume capacity
Complete this task to increase the volume capacity of a CKD host system.

Viewing array disk drive module information
Complete this task to view array disk drive module (DDM) information.

Adding arrays to ranks
Complete this task to add arrays to ranks and to specify the storage type.

Adding capacity to an extent pool (automatic mode)
Complete this task to add capacity to an existing extent pool by using the automatic configuration mode.

Adding capacity to an extent pool (manual mode)
Complete this task to add capacity to an existing extent pool by using the manual configuration mode.

Adding space-efficient storage
Complete this task to add space-efficient storage to an existing extent pool.

Modifying a single extent pool
Complete this task to modify a single extent pool that you have created.

Modifying multiple extent pools
Complete this task to modify the properties of more than one extent pool.

Removing capacity from an extent pool
Complete this task to remove capacity from one or more existing extent pools.

Modifying volume groups
Complete this task to modify volume group properties and parameters.

Initializing track space-efficient FB volumes
Complete this task to remove (initialize) the data on track space-efficient volumes for fixed block (FB) host systems. This action releases the physical extents owned by a track space-efficient volume, and leaves the volume definition in place so that the volume can be used again.

Initializing track space-efficient CKD volumes
Complete this task to remove (initialize) the data on track space-efficient volumes for CKD host systems. This action releases the physical extents owned by a track space-efficient volume, and leaves the volume definition in place so that the volume can be used again.

Modifying CKD volumes
Complete this task to modify the properties and attributes of established CKD volumes.

Modifying CKD LCUs
Complete this task to modify the attributes and properties of established CKD logical control units (LCUs).

2.1. Viewing and modifying properties

Complete this task to access properties pages so that you can review or modify the properties.

Procedure

1. In the main page of the target subject area (for example, Hosts), select an item in the table.
2. In the **Action** drop-down list, select **Properties**. The properties page for the selected item is displayed.
3. If the properties page is solely informational, the only button option is **OK**. Click it to close the page. If the properties page has attributes that you can modify, the button options are **OK** or **Apply**. Click **OK** to apply the changes and close the page. Click **Apply** to apply the changes and leave the page open.

Parent topic: Managing logical storage configurations
Parent topic: Storage complexes
Parent topic: Storage images

Related tasks
Printing a report
Downloading spreadsheet information

Related information
Host attachment

2.2. Increasing FB volume capacity

Complete this task to increase the volume capacity of a fixed block (FB) host system.

About this task

To increase the capacity of a fixed block (FB) volume, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **FB Volumes**. The FB Volumes - Main page is displayed.
2. On the FB Volumes - Main page, select a storage image from the drop-down list. In the Tasks section, click **Manage existing volumes**. The Manage volumes page is displayed.
3. Select the volumes for which you want to increase capacity. You can select single or multiple volumes.
4. In the **Action** menu, choose **Increase capacity**. The Increase capacity page is displayed.
5. Select a **Size format** that you want to display.
6. Specify a new capacity for the volumes that is between the minimum size and available capacity values displayed, and then click **OK**. The most efficient method of allocating storage is in exact multiples of 1 GB binary (one extent) or in increments of 2 million blocks.

   Note: The minimum size for the volume expansion will be the current size of the largest capacity of the selected volumes (volume size cannot be decreased, so the minimum size will be set to correspond to the largest selected volume). When inputting the new capacity value, only integer values are acceptable. For example, you cannot set the size format to DS GB and specify a value of 1.5 GB. You must use whole integers such as 1 GB or 2 GB. For finer resolution, the blocks format should be used.

7. A long running task indicator is displayed. Click **Close** to dismiss the window and return to the Manage volumes page, **End task** to stop the long running task and leave the volume capacity at its original size, or **View Summary** to view additional details about the long running task.

Parent topic: Managing logical storage configurations
Library | Support | Terms of use | Feedback

2.3. Increasing CKD volume capacity

Complete this task to increase the volume capacity of a CKD host system.

About this task

To increase the capacity of a CKD volume, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. The CKD LCUs and Volumes - Main page is displayed.
2. Select a storage image from the drop-down list. In the Tasks section, click **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
3. Select one or more LCUs with the volumes that you want to increase the capacity for. The volumes associated with those LCUs are displayed in the Manage CKD volumes table.
4. Select one or more volumes and select **Increase capacity** from the **Action** drop-down list. The Increase capacity page is displayed.
5. Select a **Size format** that you want to display.
6. Specify a new capacity for the volumes that is between the minimum size and available capacity values displayed, and then click **OK**.

   Note: The minimum size for the volume expansion will be the current size of the largest capacity of the selected volumes (volume size cannot be decreased, so the minimum size will be set to correspond to the largest selected volume). When inputting the new capacity value, only integer values are acceptable.

7. A confirmation message is displayed. Click **OK** to change the size of all of the volumes that you have selected, or click **Cancel** to cancel the operation without changing the size of the selected volumes.
8. A long running task indicator is displayed. Click **Close** to dismiss the window and return to the Manage CKD LCUs page, **End task** to stop the long running task and leave the volume capacity at its original size, or **View Summary** to view additional details about the long running task.

Parent topic: Managing logical storage configurations

2.4. Viewing array disk drive module information

Complete this task to view array disk drive module (DDM) information.

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Arrays tab, select one or more arrays. Select **DDM Properties** from the **Action** drop-down list. The DDM properties page is displayed.
3. When you are finished reviewing the DDM information, click **OK** to close the page.

Parent topic: Managing logical storage configurations

Related reference
Create array — Add arrays to rank
Single array properties — Status

2.5. Adding arrays to ranks

Complete this task to add arrays to ranks and to specify the storage type.

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Arrays tab, select an array. Select **Create arrays** from the **Action** drop-down list. The Create arrays page is displayed.
3. Select **Add Arrays to Rank**. The Create arrays - Add arrays to ranks page is displayed.
4. Accept or change the storage type.
5. Click the **OK** button to apply the array.

Parent topic: Managing logical storage configurations

Related reference
Create array — Add arrays to rank
Single array properties — Status

2.6. Adding capacity to an extent pool (automatic mode)

Complete this task to add capacity to an existing extent pool by using the automatic configuration mode.

About this task

The system automatically completes the capacity configuration process for an existing extent pool when you supply the required information and leave the Type of Configuration set to **Automatic**.

After you have supplied the required information, click **OK** to capture the capacity value changes in the Add capacity to pool verification table.

Perform the following steps to use the system to automatically configure and add capacity to the designated extent pool:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pool to which you want to add capacity. Select **Add capacity to pools** from the **Action** drop-down list. The Add capacity to extent pool page is displayed with the Type of Configuration set to **Automatic**.
3. Supply the following information to allow the system to perform an automatic capacity configuration for the extent pool:

   RAID

   Select one of the following RAID types:

   - RAID 5
   - RAID 6
   - RAID 10

   RAID type is used by the ranks in the extent pool. RAID is a methodology for grouping disk drives to protect against data loss from a failing disk drive.

   Type of Configuration

   Leave the selection set to Automatic.

   DA Pair Usage

   Select one of the following from the pull-down list:

   - Spread among all pairs
   - Spread among least used pairs
   - Sequentially fill all pairs

   The bar graph displays in real-time the results of your choice. This information is used to analyze the best configuration selection for your system.

   Drive Class

   Select one of the disk drive modules in the pull-down list. If one of the selections within this field is *146 GB 15K Enterprise Drive with xxxx GB available*, this example represents the following designation:

   - The disk drive module is a high-speed fibre channel disk drive.
   - The disk drive module has a size of 146 GB and a speed of 15 000 rpm.

   When you make your selection, an information box is displayed that lists the available capacity as ranks and the capacity size that is needed to create the extent pool.

   Select Capacity to Configure

   Select the appropriate rank designation for the capacity that you want to configure. The bar

graph displays the results of the choice that you have selected.

4. Click **OK** to add the extent pool to the Add capacity to pool verification table. You can review the information in the Add capacity to pool verification table before continuing. If you decide more capacity is required, you can select the extent pool in the table and select **Add Capacity to Pool . . .** from the table drop-down list. When you are satisfied with the capacity for the selected extent pool, click **Create All** on the Add capacity to pool verification table to initiate the process and add the specified capacity to the system. The Add capacity to pool verification table is refreshed to show the added capacity. Click **Cancel** to cancel the configuration process.

Parent topic: Managing logical storage configurations

Related tasks
Adding capacity to an extent pool (manual mode)

Related reference
Add capacity to pool
Add capacity pool verification table

2.7. Adding capacity to an extent pool (manual mode)

Complete this task to add capacity to an existing extent pool by using the manual configuration mode.

About this task

If you want to control the content of an extent pool, you can take advantage of the manual configuration mode to add capacity to the extent pool. When you change the Type of Configuration option from automatic to manual, the Add Capacity to an extent pool page changes and displays the table of available arrays, array site, and ranks. You can select from the list of available components to create your extent pool. After you make your selections, the results are displayed in real-time in the DA Pair bar graph. The system uses this information to create the extent pool.

Perform the following steps to add capacity to an extent pool using the manual process:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage – Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pool to which you want to add capacity. Select **Add capacity to pools** from the **Action** drop-down list. The Add capacity to extent pool page is displayed.
3. Observe the displayed value for RAID and determine whether you must change it.

   Use the RAID pull-down list to select one of the following values:

   - RAID 5
   - RAID 6
   - RAID 10

   RAID type is used by the ranks in the extent pool. RAID is a methodology for grouping disk drives to protect against data loss from a failing disk drive. An extent pool can contain ranks of different RAID types. An extent pool of this type is listed with a RAID type of mixed on the Extent Pools – Main page.
4. Select **Manual** from the Type of Configuration drop-down list.

   The Select Available Capacity section of the Add Capacity to Pool page changes format and displays a table of available arrays, array sites, and ranks that can be used to add capacity to the extent pool. In addition, the bar graph is retained.

   Note: If there are no values that are displayed in the table, select a different RAID type.
5. Select the components in the table that you want to use to add capacity to the extent pool. The effects of adding capacity are immediately reflected in the bar graph.
6. Click **OK** to add the extent pool to the Add capacity to pool verification table. You can review the information in the Add capacity to pool verification table before you continue. If you decide more capacity is required, you can select the extent pool in the table and select **Add Capacity to Pool . . .** from the table drop-down list. When you are satisfied with the capacity for the selected extent pool, click **Create All** on the Add capacity to pool verification table to initiate the process and add the specified capacity to the system. The Add capacity to pool verification table is refreshed to show the added capacity. Click **Cancel** to cancel the configuration process.

Parent topic: Managing logical storage configurations

Related tasks
Adding capacity to an extent pool (automatic mode)

Related reference
Add capacity to pool
Add capacity pool verification table

2.8. Adding space-efficient storage

Complete this task to add space-efficient storage to an existing extent pool.

About this task

A logical volume that allocates space only during a write transaction is referred to as a *track space-efficient* volume. When you use the storage allocation method that assigns track space-efficient volumes, you increase the potential for a more effective use of your system storage. Track space-efficient volumes are useful when you are engaged in a data mining process, testing a new application, or backing up data to tape. Track space-efficient volumes are created for FB and CKD hosts using the Create Volumes wizard.

To add space-efficient storage to one or more extent pools, perform the following steps:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pool or pools to which you want to add a space-efficient repository. Select **Add space-efficient repository** from the **Action** drop-down list. The Add space-efficient repository page is displayed with default values in each input field.

   Notes:

   a. If you select multiple extent pools and make changes, the changes are applied to all the extent pools.

   b. If any selected extent pool already has space-efficient storage, the Add space-efficient repository action is canceled and an error message is displayed.

3. Enter values as needed in the following input fields:

   Repository Capacity (GB)

   Enter the amount of real capacity that is necessary to meet the actual physical capacity that is allocated to the space-efficient storage volumes during write transactions.

   Repository Capacity Threshold (%)

   Enter the threshold percentage of allocated capacity out of the total repository capacity of the extent pool. The system generates an alert when this repository threshold percentage is reached. This allows you to make adjustments before a repository full condition occurs.

4. Click **OK** to initiate the add space-efficient storage process.

Parent topic: Managing logical storage configurations

2.9. Modifying a single extent pool

Complete this task to modify a single extent pool that you have created.

About this task

The following list shows the type of extent pool information that you can modify from within the single extent pool properties pages. The following information can be modified on the properties General tab:

- Extent pool name
- Storage threshold percentage
- Storage reserved percentage
- Repository capacity threshold percentage

Perform the following steps to modify an extent pool:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pool that you want to modify. Select **Properties** from the **Action** drop-down list. The Single Pool Properties page is displayed with the General tab opened.
3. Enter changes in the following fields. No changes occur unless you click the **OK** or **Apply** buttons.

   Extent Pool Name

   You can enter a new name for the extent pool. A valid extent pool name contains up to 16 characters including the prefix, the name, and the suffix.

   Storage Threshold (%)

   You can enter a number 0 - 100. This threshold value represents a percentage of allocated extents out of the total nonreserved storage extents. The system generates an alert when this threshold percentage is reached. This allows you the time to make adjustments before a storage full condition occurs.

   Storage Reserved (%)

   You can enter a number 0 - 100. This value represents the percentage of extents in the extent pool that are reserved and cannot be allocated for use.

4. Click **OK** or **Apply** to commit your changes.

   If you click **OK**, the change is reflected in the extent pool table listing and the properties page is closed. If you click **Apply**, the extent pool table listing is updated and the properties page does not disappear.

   Note: Clicking **OK** or **Apply** commits the changes made to both the General tab and the Space-efficient Storage tab. You can switch back and forth between the tabs, but when the changes are committed they are committed for both tabs.

5. Click the Space-efficient Storage tab and determine whether you want to make changes to the following values.

   Repository Capacity Threshold (%)

   You can enter a number 0 - 100. This value represents a percentage of the total repository capacity. If the percentage of repository capacity that is allocated to track space-efficient volumes reaches the specified percentage, the system issues an alert.

6. Click **OK** or **Apply** to commit your changes.

   If you click **OK**, the change is reflected in the extent pool table listing and the properties page is closed. If you click **Apply**, the extent pool table listing is updated and the properties page remains open, which permits more changes.

Note: Clicking **OK** or **Apply** commits the changes made to both the General tab and the Space-efficient Storage tab. You can switch back and forth between the tabs, but when the changes are committed they are committed for both tabs.

Parent topic: Managing logical storage configurations
Library | Support | Terms of use | Feedback

2.10. Modifying multiple extent pools

Complete this task to modify the properties of more than one extent pool.

About this task

If you want to modify the properties of more than one extent pool at the same time, you can do so by following this procedure. It is important to remember that any change that you make and commit to the system is transferred to all the extent pools that you are querying.

Perform the following steps to modify the properties of multiple extent pools:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pools that you want to modify. Select **Properties** from the **Action** drop-down list. The Multiple Pools Properties page is displayed.
3. Enter changes in the following fields. No changes occur unless you click the **OK** or **Apply** buttons.

   Extent Pool Name Prefix

   This field allows you to change the name prefix consisting of 1 to 12 characters that is used on all your selected extent pools.

   Prefixes Already In Use

   A text box of the extent pool prefixes that are currently in use. Use this list to ensure that you are not assigning a duplicate prefix to an extent pool.

   Storage Threshold (%)

   You can enter a number 0 - 100. This threshold value represents a percentage of allocated extents out of the total nonreserved storage extents. The system generates an alert when this threshold percentage is reached. This allows you the time to make adjustments before a storage full condition occurs.

   Storage Reserved %

   You can change the percentage of the total extent pool capacity that is reserved. This percentage is prevented from being allocated to volumes or space-efficient storage. You can specify a value of 0 - 100, with 0 being the default. If you specify 0, no extents are reserved.

   Repository Capacity Threshold (%)

   You can change the threshold percentage of allocated capacity out of the total repository capacity of the extent pool. The system generates an alert when the repository threshold percentage is reached. This allows you to make adjustments before a repository full condition occurs. You can specify a value of 0 - 100 with 100 being the default.

4. Click **OK** or **Apply** to commit your changes.

   If you click **OK**, the change is reflected in the extent pool table listing and the properties page is closed. If you click **Apply**, the extent pool table listing is updated and the properties page remains open, which permits more changes.

Parent topic: Managing logical storage configurations

2.11. Removing capacity from an extent pool

Complete this task to remove capacity from one or more existing extent pools.

About this task

You can reduce the capacity of an extent pool by removing ranks from the extent pool. After you select the ranks that you want to remove and initiate the removal process, a confirmation message is displayed. You are asked to confirm the removal of any volumes that might be associated with the ranks that you want removed.

Attention: When you confirm that you want the remove capacity process to complete, all the associated data is lost. There is no way to recover the data, unless backups were made of this data.

To reduce the capacity of an extent pool, perform the following steps:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. Select a storage image from the drop-down list.
2. On the Extent pools tab, select the extent pools from which you want to remove capacity. Select **Remove capacity from pools...** from the **Action** drop-down list. The Remove capacity from extent pools page is displayed. This page contains a table that lists all the ranks associated with the selected extent pools for the storage image that you are accessing.
3. Select one or more of the ranks that you want removed to reduce the capacity of the extent pools.
4. Click **Remove** to initiate the deletion process. A confirmation message is displayed. If there are volumes that are associated with the ranks that are being deleted, the confirmation message contains a list of the volumes.
5. Verify that you want to delete the ranks and volumes that are listed in the confirmation message, and click **Continue** to complete the deletion process.

Parent topic: Managing logical storage configurations

Related reference
Remove capacity from an extent pool

2.12. Modifying volume groups

Complete this task to modify volume group properties and parameters.

Before you begin

The Properties option is available in the **Action** drop-down list only after you have selected one of the groups in the table.

Procedure

1. In the navigation, select **Volumes** > **Volume groups**. On the Volume groups main page, in the **Storage image** drop-down list, select the storage image that contains the volume group that you want to modify.
2. Select the volume group that you want to modify. In the **Action** drop-down list, select **Properties**. The Volume Group Properties page is displayed.
3. On the Volume Group Properties page you can change the nickname of the volume group, or you can select volumes to add or remove from the volume group. The table displays all of the volumes that are assigned to this volume group.
   a. To remove one or more volumes from the volume group, select the volumes on the table of assigned volumes and select **Remove volumes** from the **Action** drop-down menu. A warning message is displayed.
   b. To add one or more volumes to the volume group, select **Add volumes** from the **Action** drop-down menu. The Volume groups page is displayed. Select the volumes that you want to add to the volume group. Click **Next** and follow the Volume groups wizard to complete the process.
4. Verify that the properties are correct, and click **OK** to save your changes or click **Cancel** to exit without saving your changes.

Parent topic: Managing logical storage configurations

Related concepts
Open systems

Related tasks
Initializing track space-efficient FB volumes

2.13. Initializing track space-efficient FB volumes

Complete this task to remove (initialize) the data on track space-efficient volumes for fixed block (FB) host systems. This action releases the physical extents owned by a track space-efficient volume, and leaves the volume definition in place so that the volume can be used again.

Procedure

1. In the navigation, select **Volumes** > **FB Volumes**. The FB Volumes - Main page is displayed.
2. Select a storage image. In the Tasks box, click **Manage existing volumes**. The Manage volumes page is displayed.
3. Select the volume or volumes that you want to initialize.

   Note: Only track space-efficient volumes can be initialized.

4. From the **Action** drop-down list, select **Initialize volumes**. The Initialize volumes confirmation page is displayed. You can review the volumes that you selected to initialize before deleting the data.
5. Select **Continue** to remove all data from the volumes and initialize, while leaving the volumes intact, or click **Cancel** to exit without initializing the volumes.

Parent topic: Managing logical storage configurations

Related tasks
Modifying volume groups
Modifying CKD volumes

Library | Support | Terms of use | Feedback

2.14. Initializing track space-efficient CKD volumes

Complete this task to remove (initialize) the data on track space-efficient volumes for CKD host systems. This action releases the physical extents owned by a track space-efficient volume, and leaves the volume definition in place so that the volume can be used again.

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. The CKD LCUs and Volumes - Main page is displayed.
2. Select a storage image. In the Tasks box, click **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
3. Select an LCU from the table. The Manage Volumes table is displayed, with the volumes associated with the selected LCUs.
4. Select the volume or volumes that you want to initialize.

   Note: Only track space-efficient volumes can be initialized.
5. From the **Action** drop-down list, select **Initialize volumes**. The Initialize volumes confirmation page is displayed. You can review the volumes that you selected to initialize before deleting the data.
6. Select **Continue** to remove all data from the volumes and initialize, while leaving the volumes intact, or click **Cancel** to exit without initializing the volumes.

Parent topic: Managing logical storage configurations
Library | Support | Terms of use | Feedback

2.15. Modifying CKD volumes

Complete this task to modify the properties and attributes of established CKD volumes.

About this task

To modify CKD volumes, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. On the CKD LCUs and Volumes – Main page, select the storage image that includes the volume that you want to modify. In the Tasks box, select **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
2. Select an LCU from the table. The Manage Volumes table is displayed, with the volumes associated with the selected LCUs.
3. Select the volume or volumes that you want to modify. In the **Action** drop-down list, select **Properties**. If you selected one volume, the Single volume properties page is displayed, with the General tab open. If you selected more than one volume, the Multiple Volumes Properties page is displayed.
4. You can modify the volume nickname and the performance group.
5. Click **OK** to apply the changes and close the page. Click **Apply** to apply the changes and leave the page open.

Parent topic: Managing logical storage configurations

Related tasks
Initializing track space-efficient FB volumes
Modifying CKD LCUs

Library | Support | Terms of use | Feedback

2.16. Modifying CKD LCUs

Complete this task to modify the attributes and properties of established CKD logical control units (LCUs).

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. The CKD LCUs and Volumes - Main page is displayed. In the **Select storage image** list, select the storage image that includes the LCUs that you want to modify.
2. In the Tasks box, click **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
3. Select the LCU or LCUs that you want to modify. In the **Action** list, select **Properties**. If you selected one volume, the Single LCU properties page is displayed, with the General tab open. If you selected more than one volume, the Multiple LCUs properties page is displayed.
4. Define the parameters to modify the selected LCU:
    a. Specify the LCU type.
    b. Change the default SSID by highlighting and typing over it.
    c. Enable or disable critical mode.
    d. Establish the timeout times in seconds.
    e. Enable or disable a consistency group, and, as necessary, establish the long busy timeout in seconds.
5. Once you have modified the values, click **OK** to complete the process and close the page or click **Apply** to apply the changes and leave the page open.

Parent topic: Managing logical storage configurations

Related concepts
Logical control unit

Related tasks
Modifying CKD volumes

3. Deleting a logical storage configuration

Follow these steps to delete or remove a logical storage configuration for fixed block (FB) storage images. This applies to entire configurations including, ranks, arrays, volumes, and extent pools.

About this task

Recommendation: Because this procedure deletes the host and logical storage configuration from the applicable storage unit you can, prior to performing this task, export and save the current configuration so that it can be restored.

When you are deleting a logical storage configuration, individually delete the hosts and volume group.

To delete a logical configuration, perform the following steps:

Procedure

1. Delete host access to the volume groups that you want to remove. To remove the host port IDs use the subtask in Deleting hosts.
2. Delete the volume group. To delete volume groups, you must complete the subtask in Deleting FB volume groups.
3. Delete the array. To delete an array, you must complete the subtask in Deleting arrays.

   Note: Any ranks that make up an extent pool, along with the volumes within the extent pool, are automatically deleted if the rank is assigned to the array when you delete the array.
4. Delete any unassigned ranks using the subtask in Deleting ranks.
5. Delete the extent pools, if they were not deleted by prior delete actions. To delete the extent pools for FB volumes, you must complete the subtask in Deleting extent pools.

What to do next

Verify that no logical configuration exists for the applicable storage image.

Deleting hosts
Complete this task to delete an existing host.

Deleting FB volume groups
Complete this task to delete fixed block (FB) volume groups, releasing any assigned volumes for assignment to other volume groups.

Deleting arrays
Complete this task to delete arrays with and without volumes using storage.

Deleting ranks
Complete this task to delete ranks with or without volumes using storage.

Deleting FB volumes
Complete this task to delete fixed block (FB) volumes from the storage image.

Deleting extent pools
Complete this task to delete the selected extent pools.

Deleting space-efficient storage
Complete this task to delete space-efficient storage.

Deleting CKD volumes
Complete this task to delete CKD volumes, and any of their associated aliases that are not aliases for other volumes.

Deleting CKD LCUs
Complete this task to delete CKD logical control units (LCUs) and any of their associated volumes.

Related concepts
Deleting data storage configurations

Related tasks

Deleting a fixed block data storage configuration
Deleting a count key data storage configuration

3.1. Deleting hosts

Complete this task to delete an existing host.

About this task

Perform the following steps to delete an existing host:

Procedure

1. In the navigation, select **Hosts** > **Hosts**. Select a storage image. The Hosts — Main page is displayed.
2. In the Hosts — Main page, select **Manage hosts** from the **Tasks** menu. The Manage Hosts page is displayed.
3. In the Manage Host page, select the host that you want to delete.
4. From the **Action** menu, select **Delete**. A confirmation window indicates that you want to delete the host and the associated host ports. If the hosts are mapped to a volume group, you might lose access to the volume group.
5. Click **Continue** to delete the host, or click **Cancel** to cancel the operation without deleting the host.

Parent topic: Deleting a logical storage configuration
Parent topic: Hosts

Related concepts
Host attachment overview

Related tasks
Creating new hosts
Managing hosts

Related information
Host attachment

3.2. Deleting FB volume groups

Complete this task to delete fixed block (FB) volume groups, releasing any assigned volumes for assignment to other volume groups.

About this task

If the number of volumes in this volume group is not zero, the volumes that are assigned to this volume group go into an unassigned mode where they are eligible for assignment to other volume groups.

Note: The volume group is the link between the FB volumes and the hosts. If you delete a volume group that is attached to a host identifier, the host is not able to access the volumes.

Procedure

1. In the navigation, select **Volumes** > **Volume groups**. On the Volume groups main page, in the **Select storage image** drop-down list, select the storage image that contains the volume group that you want to delete.
2. Select the volume group or groups that you want to delete.
3. In the **Action** drop-down list, select **Delete**. A confirmation dialog box is displayed.
4. Confirm the action to complete the deletion of the selected volume groups.

Parent topic: Deleting a logical storage configuration

Related concepts
Open systems

3.3. Deleting arrays

Complete this task to delete arrays with and without volumes using storage.

Before you begin

If the arrays you are deleting are assigned to volume groups, you must first remove the arrays from the volume groups before you can delete them.

If the selected arrays have volumes using storage on the array, an **Attention** message indicates the volumes that have storage configured on the arrays. The message also indicates that the volumes will be deleted as part of the array deletion process.

Procedure

1. **Pools** > **Internal Storage** > **Arrays** > **Delete Arrays** > **OK**.
2. Select the array or arrays that you want to delete. In the **Action** drop-down menu, select **Delete Arrays**.

   Note: If the array that you delete has volumes using storage on it, these volumes will be deleted along with the array. You will see a message displaying the volume numbers and nicknames that will be deleted and be asked whether you want to continue with or cancel the deletion process.

   A confirmation message is displayed. If the volumes are assigned, regardless of any storage used, the confirmation indicates that the array will be deleted.
3. Click **OK** to complete the deletion of the selected arrays.

Parent topic: Deleting a logical storage configuration

Related reference
Create array — Add arrays to rank
Single array properties — Status

3.4. Deleting ranks

Complete this task to delete ranks with or without volumes using storage.

Procedure

1. In the navigation, select **Pools** > **Internal Storage**. The Internal storage - Main page is displayed. In the **Select storage image** drop-down list, select the storage image that contains the rank that you want to delete.
2. On the Ranks tab, select the rank or ranks that you want to delete. In the **Action** drop-down list, select **Delete**. A confirmation dialog box is displayed.

   Note: If the ranks you selected have volumes using storage on the rank, an **Attention** message lists the volumes that have storage configured on the ranks. The message also indicates that the volumes will be deleted as part of the rank deletion process.
3. If the **Attention** message is displayed, click **Continue** or **Cancel**. If you click **Cancel**, you are returned the Internal storage - Main page. If you click **Continue**, go to the next step.
4. After you click **Continue**, a confirmation dialog is displayed. If the selected ranks are unassigned or do not have any extents allocated on the ranks, confirming the dialog box deletes the selected ranks. The array assigned to this rank goes into an unavailable status. In an unassigned status, the array becomes eligible for assignment to another rank.
5. Confirm the dialog box to complete the deletion of the selected ranks and arrays.

Parent topic: Deleting a logical storage configuration
Library | Support | Terms of use | Feedback

3.5. Deleting FB volumes

Complete this task to delete fixed block (FB) volumes from the storage image.

Procedure

1. In the navigation, select **Volumes** > **FB Volumes**. The FB Volumes – Main page is displayed. In the **Select storage image** drop-down list, select the storage image that contains the volume that you want to delete. In the Tasks box, click **Manage existing volumes**. The Manage volumes page is displayed.
2. Select the volume or volumes that you want to delete.
3. In the **Action** drop-down list, select **Delete**. A confirmation dialog box is displayed.
4. Confirm the action to complete the deletion of the selected volumes.

Parent topic: Deleting a logical storage configuration

Related concepts
Open systems

3.6. Deleting extent pools

Complete this task to delete the selected extent pools.

About this task

If you want to delete extent pools, it is important to remember that all of the data, space-efficient storage, and the associated volumes are also deleted. It is not possible to recover the deleted items.

Perform the following steps to delete extent pools:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. In the **Select storage image** drop-down list, select the storage image that contains the extent pools that you want to delete.
2. On the Extent pools tab, select the extent pools that you want to delete.
3. Select **Delete Pools . . .** from the table drop-down list. The Delete Extent Pools warning page is displayed. It shows which extent pools will be deleted and a list of associated volumes that will also be deleted.
4. Verify that you want to delete the selected extent pools by clicking **Continue**.

Parent topic: Deleting a logical storage configuration
Library | Support | Terms of use | Feedback

3.7. Deleting space-efficient storage

Complete this task to delete space-efficient storage.

About this task

If you want to delete space efficient storage, it is important to remember that all of the data and the associated volumes are deleted and unrecoverable.

Perform the following steps to delete space efficient storage:

Procedure

1. In the navigation, select **Pools** > **Internal storage**. The Internal storage - Main page is displayed. In the **Select storage image** drop-down list, select the storage image that contains the space efficient storage that you want to delete.
2. On the Extent pools tab, select the extent pools for which you want to delete space efficient storage.
3. Select **Delete Space Efficient Storage . . .** from the table drop-down list. The Delete Space Efficient Storage warning page is displayed. It shows which extent pools and associated volumes will be deleted.

   Note: The data that is associated with these volumes is destroyed during the deletion process and is not recoverable after the deletion process completes.
4. Verify that you want to delete the space-efficient storage by clicking **Continue**.

Parent topic: Deleting a logical storage configuration
Library | Support | Terms of use | Feedback

3.8. Deleting CKD volumes

Complete this task to delete CKD volumes, and any of their associated aliases that are not aliases for other volumes.

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. The CKD LCUs and Volumes - Main page is displayed. In the **Select storage image** list, select the storage image that includes the LCUs that you want to modify.
2. In the Tasks box, click **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
3. Select the LCU that includes the volume or volumes that you want to delete. The Manage Volumes table is displayed, with the volumes associated with the selected LCU.
4. Select the volumes that you want to delete. This process deletes the selected volumes and any associated aliases that are not aliases for other volumes.
5. In the **Action** list, select **Delete**. A confirmation window is displayed.
6. Confirm the action to complete the process.

Parent topic: Deleting a logical storage configuration

3.9. Deleting CKD LCUs

Complete this task to delete CKD logical control units (LCUs) and any of their associated volumes.

Before you begin

If the selected LCUs have volumes defined in them, an **Attention** message indicates that there are volumes configured in the LCUs. The message also indicates that these volumes are deleted as part of the LCU deletion process.

About this task

To delete CKD LCUs, perform the following steps:

Procedure

1. In the navigation, select **Volumes** > **CKD LCUs and Volumes**. The CKD LCUs and Volumes – Main page is displayed. In the **Select storage image** list, select the storage image that includes the LCUs that you want to modify.
2. In the Tasks box, click **Manage existing LCUs and volumes**. The Manage LCUs page is displayed.
3. Select the LCU or LCUs that you want to delete. In the **Action** list, select **Delete**. A confirmation window is displayed.
4. If the **Attention** message is displayed, click **Continue** or **Cancel**. Clicking **Cancel** returns you to the LCUs page. If you click **Continue**, go to the next step.
5. After you click **Continue**, a confirmation window is displayed and, if confirmed, the selected LCUs and volumes are deleted. If the selected LCUs do not have volumes configured in them, a confirmation window is invoked and, if confirmed, the characteristics of the selected LCUs are deleted.
6. Confirm to complete the deletion of the selected LCUs and volumes.

Parent topic: Deleting a logical storage configuration
Library | Support | Terms of use | Feedback

4. Activating licensed functions

After the IBM® service representatives have completed your DS8000® storage complex installation, your first step is to activate your licensed functions.

To activate your licensed functions, you must perform the following actions:

- Obtain your feature activation codes.
- Apply the activation codes to your storage unit. You can apply the activation codes by importing a file that you download from the IBM Disk Storage Feature Activation (DSFA) website.

The initial enablement of any optional DS8000 licensed function is a concurrent activity (assuming that the appropriate level of microcode is installed on the machine for the given function).

Note: The following activities are non-disruptive, but take effect at the next machine IML.

- Removal of a DS8000 licensed function to deactivate the function.
- A lateral change or reduction in the license scope. A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.

Obtaining activation codes
Complete this task to obtain your feature activation codes for the licensed features for each storage unit. To obtain your activation codes, you must connect to the IBM Disk Storage Feature Activation (DSFA) website.

Adding activation keys
Complete this task to apply the activation codes that enable you to begin configuring storage on a storage image.

Importing activation keys
Complete this task to import the activation keys that must be applied before you can begin configuring storage on a storage unit.

Obtaining activation codes
Complete this task to obtain your feature activation codes for the licensed features for each storage unit. To obtain your activation codes, you must connect to the IBM Disk Storage Feature Activation (DSFA) website.

Adding activation keys
Complete this task to apply the activation codes that enable you to begin configuring storage on a storage image.

Importing activation keys
Complete this task to import the activation keys that must be applied before you can begin configuring storage on a storage unit.

Parent topic: Setting up the DS Storage Manager

Related concepts
Licensing on the DS8000

Related tasks
Activating your machine and feature licenses using the DS CLI
Obtaining activation codes
Adding activation keys
Importing activation keys

4.1. Obtaining activation codes

Complete this task to obtain your feature activation codes for the licensed features for each storage unit. To obtain your activation codes, you must connect to the IBM Disk Storage Feature Activation (DSFA) website.

Before you begin

Before you can connect to the site, ensure that you have the following items:

- The IBM License Function Authorization documents. If you are activating codes for a new storage unit, these documents are included in the shipment of the storage unit. If you are activating codes for an existing storage unit, IBM sends these documents to you in an envelope.
- A removable media for downloading your activation codes into a file. Use the removable media if you cannot access the IBM® System Storage® DS Storage Manager from the system that you are using to access the DSFA website. Instead of using removable media, you can also write down the activation codes and then manually enter them into the system that runs the DS Storage Manager.
- The machine serial number, model, and signature.

About this task

To obtain your activation codes, perform the following steps:

Procedure

1. At a computer with an Internet connection and a browser, connect to the IBM Disk Storage Feature Activation (DSFA) website at IBM Storage DSFA.
2. The DSFA application displays in the browser. Use the application to obtain the activation codes and follow the instructions on the screen.

   Note: In most situations, the DSFA application can locate your 239x license authorization record when you enter the DS8000® 242x machine type, serial number and signature. However, if the 239x license authorization record is not attached to the 242x record, you must assign it to the 242x record in the DSFA application. In this situation, you will need the 239x serial number (which you can find on the License Function Authorization document).

   After you complete these steps and your license codes are not present or incorrect for your storage unit, contact your IBM representative.

Next topic: Adding activation keys
Parent topic: Activating licensed functions
Parent topic: Activating licensed functions

Related concepts
Activating licensed functions
Licensing on the DS8000

Related tasks
Activating your machine and feature licenses using the DS CLI

Related reference
Add activation key
Import key file

4.2. Adding activation keys

Complete this task to apply the activation codes that enable you to begin configuring storage on a storage image.

Before you begin

Notes:

1. The initial enablement of any optional DS8000® licensed function is a concurrent activity (assuming that the appropriate level of microcode is installed on the machine for the given function). The following activating activities are non-disruptive, but take effect at the next machine IML for Model 941 and Model 951:
   - Removal of a DS8000 licensed function to deactivate the function. Contact your IBM service representative to perform this operation.
   - A lateral change or reduction in the license scope. A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.
2. Before you begin this task, you must check the error log and resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems. See Troubleshooting for more information on troubleshooting.
3. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS8000 communication.

About this task

The easiest way to apply the feature activation keys is to download the activation keys from the IBM Disk Storage Feature Activation (DSFA) website to your local computer and then to import the file into the DS Storage Manager. If you cannot access the DS Storage Manager from the same computer that you used to access the DSFA website, you can download the file to a CD or USB flash drive, or write down the information. If you are using either of these latter methods, ensure that you have your CD or USB flash drive containing the downloaded activation keys file or your paper that contains the written activation codes before you begin the following steps.

Procedure

1. Ensure that the Import key file page is not open. You cannot have both the Add activation key page and the Import key file page open at the same time. You must close one to access the other.
2. In the navigation, select **Home** > **System Status**. On the System Status main page, select the storage image for which you want to import the activation keys.
3. From the **Action** menu, select **Storage Image** > **Add Activation Key**. The Add Activation Key page is displayed.
   a. If you already imported your activation keys from a file or retrieved existing codes from the storage unit, the values are displayed in the fields and you can modify or overwrite them, as appropriate.
   b. If you have downloaded the activation key from the Disk Storage Feature Activation (DSFA) website, select **Import Key File** from the **Action** list and follow the prompts.
   c. If you did not download your activation key, select **Add Activation Key** and follow the prompts.
4. Click **OK** to complete the process.

Previous topic: Obtaining activation codes
Next topic: Importing activation keys
Parent topic: Activating licensed functions
Parent topic: Activating licensed functions
Parent topic: Storage images

Related concepts
Activating licensed functions

Licensing on the DS8000

Related tasks
Activating your machine and feature licenses using the DS CLI

Related reference
Add activation key
Import key file

4.3. Importing activation keys

Complete this task to import the activation keys that must be applied before you can begin configuring storage on a storage unit.

Before you begin

Notes:

1. The initial enablement of any optional DS8000® licensed function is a concurrent activity (assuming the appropriate level of microcode is installed on the machine for the given function). The following activating activities are non-disruptive, but take effect at the next machine IML:
   - Removal of a DS8000 licensed function to deactivate the function.
   - A lateral change or reduction in the license scope. A lateral change is defined as changing the license scope from fixed block (FB) to count key data (CKD) or from CKD to FB. A reduction is defined as changing the license scope from all physical capacity (ALL) to only FB or only CKD capacity.
2. Before you begin this task, you must resolve any current DS8000 problems. Contact IBM Support for assistance in resolving these problems.
3. Before you configure, disable or provide paths through any firewalls, because they might interfere with DS8000 communication.

Perform the following steps to import your activation codes:

Procedure

1. In the navigation, select **Home** > **System Status**. On the System Status main page, select the storage image for which you want to import the activation key.
2. From the **Action** menu, select **Storage Image** > **Add Activation Key**. The Add Activation Key page is displayed.
3. Click **Import key file**. The Import page is displayed.
4. In the **Select file to import** field, specify the target file. Use the **Browse** button to navigate to the appropriate directory.
5. After you have specified the key file, click **Next** to complete the process.

Previous topic: Adding activation keys
Parent topic: Activating licensed functions
Parent topic: Activating licensed functions
Parent topic: Storage images

Related concepts
Activating licensed functions
Licensing on the DS8000

Related tasks
Activating your machine and feature licenses using the DS CLI

Related reference
Add activation key
Import key file

5. Attaching hosts

This section provides general information about DS8000® host attachment and describes the specific hosts that you can attach to the DS8000 storage unit.

Apple Macintosh host attachment
This section provides the requirements that you need to verify before you attach an Apple Macintosh host system to a DS8000 with Fibre Channel adapters.

Fujitsu PrimePower host attachment

HP Tru64 UNIX host attachment
This section describes how you attach a Hewlett-Packard (HP) AlphaServer running the Tru64 UNIX operating system to a DS8000 with Fibre Channel adapters.

HP OpenVMS host attachment
This section describes how you attach a Hewlett-Packard (HP) AlphaServer or Integrity server running the OpenVMS operating system to a DS8000 with Fibre Channel adapters.

HP-UX host attachment
This section describes how you attach a Hewlett-Packard (HP) server running the HP-UX operating system to a DS8000 with Fibre Channel adapters.

IBM Power Systems hosts running IBM i
This section describes how you attach an IBM® i host system to a DS8000 with Fibre Channel adapters.

IBM Power Systems hosts running AIX host attachment
This section describes how you attach an IBM Power Systems™ running AIX host system to a DS8000 with Fibre Channel adapters.

IBM System z or IBM S/390 host attachment
This section describes how you attach a System z® or S/390® host to a DS8000 with a Fibre Channel, ESCON®, or FICON® adapter.

IBM SAN File System host attachment
This section describes the requirements that you need to verify before you attach a SAN File System metadata server to a DS8000 with Fibre Channel adapters.

IBM SAN Volume Controller host attachment
This section provides the requirements that you need to verify before you attach an IBM SAN Volume Controller host system to a DS8000 with Fibre Channel adapters.

IBM NAS Gateway 500 host attachment
This section describes how you attach an IBM NAS Gateway 500 host system to a DS8000 with Fibre Channel adapters.

Linux (x86 or x86_64) host attachment
This section describes how you attach a Linux host to a DS8000 with Fibre Channel adapters.

VMware ESX Server host attachment
This section describes how you attach a VMware ESX Server host to a DS8000 with Fibre Channel adapters.

Microsoft Windows Server host attachment
This section describes how you attach a Windows Server 2003 or 2008 host to a DS8000 with Fibre Channel adapters.

Novell NetWare host attachment
This section describes how you attach a Novell NetWare host system to a DS8000 with Fibre Channel adapters.

Silicon Graphics host attachment
This section describes how you attach a Silicon Graphics (SGI) host system to a DS8000 with Fibre Channel adapters.

Sun host attachment
This section describes how you attach a Sun host system to a DS8000 with Fibre Channel adapters.

Host attachment overview
The DS8000 series provides various host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and System z.

General requirements for attaching a host
Before you attach a host to the DS8000, review this list of general requirements for all hosts. Then, review the specific host requirements described in the section for each host.

Downloading and installing a host bus adapter driver
Complete this task to download and install host adapter drivers to use with the DS8000.

Host attachment general considerations
This section contains general conceptual overviews on key topics for host attachment.

IBM Power Systems host attachment
This section describes how you attach an IBM Power Systems host system to a DS8000 with Fibre Channel adapters.

Linux (x86 or x86_64) host attachment
This section describes how you attach a Linux host to a DS8000 with Fibre Channel adapters.

5.1. Host attachment overview

The DS8000® series provides various host attachments so that you can consolidate storage capacity and workloads for open-systems hosts and System z®.

The DS8000 series provides extensive connectivity using Fibre Channel adapters across a broad range of server environments.

Parent topic: Host attachment
Parent topic: Attaching hosts

Related tasks
Viewing and modifying properties
Creating new hosts
Managing hosts

5.2. General requirements for attaching a host

Before you attach a host to the DS8000®, review this list of general requirements for all hosts. Then, review the specific host requirements described in the section for each host.

Perform the following steps before you attach any host system to a DS8000:

1. Go to the *System Storage® Interoperation Center* (SSIC) website at IBM Storage SSIC for the most current information on supported hosts, operating systems, adapters, and switches.
2. Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.
3. Ensure that you can reference the following documentation:
   - The IBM® System Storage DS8000 Information Center at DS8000 Information Center.
   - The IBM Systems Information Centers at publib.boulder.ibm.com/eserver/.
4. Have an IBM service support representative install the storage unit.
5. Check the LUN limitations for your host system and verify that there are enough adapters that are installed in the server to manage the total LUNs that you want to attach. See LUN considerations for Fibre Channel attachment.
6. Review the "I/O adapter features" section of the *IBM System Storage DS8000 Introduction and Planning Guide* for detailed host adapter configuration rules and for information on host adapter and cable feature codes, host adapter balancing for optimal performance, and supported cabling distances.
7. Use the web-based DS Storage Manager or the DS CLI to define the host and I/O port configurations. Make sure that you define the worldwide port names for Fibre Channel ports. See Host configuration using the DS8000 interfaces for more information.
8. Install the adapter driver with your host adapter or use the steps defined in Downloading and installing a host bus adapter driver to download and install an updated adapter driver.
9. For further host attachment details, refer to the section dedicated to the host OS.

Parent topic: Introduction to host attachment
Parent topic: Attaching hosts

Related concepts
Host attachment general considerations

Related information
Apple Macintosh host attachment
Fujitsu PrimePower host attachment
HP OpenVMS host attachment
HP Tru64 UNIX host attachment
HP-UX host attachment
IBM NAS Gateway 500 host attachment
IBM SAN Volume Controller host attachment
IBM Power Systems host attachment
IBM System z or IBM S/390 host attachment
Linux (x86 or x86_64) host attachment
Microsoft Windows Server host attachment
VMware ESX Server host attachment
Novell NetWare host attachment
Silicon Graphics host attachment
Sun host attachment

5.3. Downloading and installing a host bus adapter driver

Complete this task to download and install host adapter drivers to use with the DS8000®.

About this task

To download and install a host adapter driver, perform the following steps:

Notes:

1. For Linux hosts, use the HBA driver provided in the distribution.
2. You must download the relevant vendor documentation for the driver that you select to correctly install and configure the host adapter.

Procedure

1. Go to the *System Storage® Interoperation Center* (SSIC) website at IBM Storage SSIC and verify that the host adapter you want to use is compatible with the DS8000 series, your host, and your host operating system.
2. Obtain a list of supported host bus adapters (HBAs), firmware, and device driver information for your host system and host adapter on the *System Storage Interoperation Center* (SSIC) website at IBM Storage SSIC.
3. If required, download the adapter driver from the adapter vendor website. Table 1 provides vendor web addresses and specific download steps. For the most current steps, see the vendor website download page.

Table 1. Host adapter driver download steps by vendor

Host adapter vendor	Web address	Steps to locate download page
AMCC/JNI	www.amcc.com	a. Click **Downloads** → **FC HBA/OEM Files**. b. Click **IBM**. c. Locate your adapter. d. Download the appropriate files for your adapter.
Emulex	www.emulex.com/downloads/ibm/oneconnect-software-kits.html	a. Click the tab for the adapter type that you are looking for. b. Locate your adapter. c. Download the appropriate files for your adapter.
Hewlett-Packard	www.hp.com	a. Enter the name of the host adapter in the **Search** field at the bottom of the Hewlett-Packard home page. b. In the search results, click the link for the driver for your operating system. c. Click **Download** to download the adapter driver. d. Return to the search results and review the links to documentation, such as installation requirements and release notes.

Host adapter vendor	Web address	Steps to locate download page
QLogic	support.qlogic.com/support/oem_ibm.asp	a. Click the DS8000 link. b. Locate your adapter. c. Download the appropriate files for your adapter.
Sun	www.sun.com/storage/san	If you are using Solaris 8 or 9, perform the following steps. If you are using Solaris 10, you can skip these steps because the SAN software is included in the operating system. a. Scroll down to the **Get the Software** section. b. Locate the current driver in the list and click the appropriate link. c. Type your **Username** and **Password** in the fields, and click **Log in and Continue**. If you do not have a user name and password, complete the registration process and return to this page when you have received them. **Note:** Sun only grants user names and passwords to customers who have purchased maintenance contracts. d. Click **Accept** to agree to the license agreement (required). e. Download the appropriate files for your adapter.

4. Follow the installation instructions from the vendor to install the host adapter driver. Some installation instructions might be included in the readme files that are included in the compressed download, but some installation documentation might have to be downloaded separately. Return to the vendor web address that is listed in Table 1 to locate installation and configuration instructions.
5. Locate host attachment information for your host in the "Attaching hosts" section of the DS8000 Information Center, and follow any additional driver configuration instructions.

Parent topic: Introduction to host attachment
Parent topic: Attaching hosts
Library | Support | Terms of use | Feedback

5.4. Host attachment general considerations

This section contains general conceptual overviews on key topics for host attachment.

Multipathing
Installing a multipathing driver on the host is highly recommended as it provides redundancy for failed paths and failed storage controllers. The DS8000 supports a variety of multipathing drivers, depending upon the OS. Refer to the System Storage Interoperation Center (SSIC) website at www.ibm.com/systems/support/storage/config/ssic/ for the correct multipathing driver.

Attachment restrictions for Copy Services
This section describes attachment restrictions for the standard FlashCopy®, FlashCopy SE, and Remote Mirror and Copy features. The most important difference between the standard FlashCopy and FlashCopy SE features is space efficiency.

Host attachment path considerations for a storage image
Path considerations affect the performance and availability of a storage image.

Fibre Channel host attachment
This section provides an overview of Fibre Channel attachment to a DS8000, including architecture, topology, and LUN access modes.

Parent topic: Attaching hosts

Related concepts
Host attachment path considerations for a storage image

Related information
General requirements for attaching a host

5.4.1. Multipathing

Installing a multipathing driver on the host is highly recommended as it provides redundancy for failed paths and failed storage controllers. The DS8000 supports a variety of multipathing drivers, depending upon the OS. Refer to the System Storage Interoperation Center (SSIC) website at www.ibm.com/systems/support/storage/config/ssic/ for the correct multipathing driver.

IBM System Storage Multipath Subsystem Device Driver
The IBM® System Storage® Multipath Subsystem Device Driver (SDD) is IBM's multipathing driver and is used with several operating systems.

Parent topic: Host attachment general considerations

5.4.1.1. IBM System Storage Multipath Subsystem Device Driver

The IBM® System Storage® Multipath Subsystem Device Driver (SDD) is IBM's multipathing driver and is used with several operating systems.

The SDD resides in the host server with the native disk device driver and takes advantage of redundant storage configurations to provide single-point failure protection, high data availability, and storage concurrent firmware download. It also provides a dynamic I/O load-balancing algorithm to reach the best I/O performance.

The SDD provides the following functions:

- Enhanced data availability
- Automatic path failover and failback
- Dynamic I/O load-balancing across multiple paths
- Multiple path selection policies for the host system
- Concurrent download of licensed machine code
- Multiple path SAN boot on storage devices that support multipath I/O (MPIO)

The term *SDD* represents both the SDD as the established multipath subsystem device driver and as the MPIO path control module, depending upon the operating system. Table 1 provides examples.

Table 1. Examples of SDD MPIO path control modules

Operating system	Description	Acronym
AIX®	SDD Path Control Module	SDDPCM
Windows	SDD Device Specific Module	SDDDSM

As new operating systems are added, the SDD will be named according to the operating system.

For more information about the SDD, see IBM Storage Software SDD.

Parent topic: Introduction to host attachment
Parent topic: Multipathing

5.4.2. Attachment restrictions for Copy Services

This section describes attachment restrictions for the standard FlashCopy®, FlashCopy SE, and Remote Mirror and Copy features. The most important difference between the standard FlashCopy and FlashCopy SE features is space efficiency.

When you copy a source volume to a target volume with the FlashCopy or Remote Mirror and Copy feature, the source and target volumes must be on different host systems to enable concurrent read/write access of both volumes.

A copy operation with the source and target volume on the same host system creates a target volume with the same identification as the source volume. The result is that the host system sees two identical volumes. When the copy operation creates the same identification for the target volume as for the source volume, you are not able to distinguish one from the other. Therefore, you might not be able to access the original data.

Novell NetWare restriction: You cannot create a host target on a single Novell NetWare host system. For Novell NetWare, the target volume must be attached to a second Novell NetWare host system.

The target volume and the source volume can be on the same host system for a FlashCopy or Remote Mirror and Copy operation only under the following conditions:

- For AIX® and Sun, when the host system is not using a logical volume manager (LVM).
- For AIX, when the host system is using an LVM with `recreatevg` command support.
- For HP using a Fibre Channel connection, when an LVM accesses a host through the `vfchigid -f` command.
- For any host system, when the host system can distinguish between a source and a target volume that have the same identification.

Parent topic: Host attachment
Parent topic: Host attachment general considerations
Library | Support | Terms of use | Feedback

5.4.3. Host attachment path considerations for a storage image

Path considerations affect the performance and availability of a storage image.

For optimal performance, use the following guidelines if you are attaching multiple paths from a single host system to I/O ports on a host adapter of a storage image:

- Use attached I/O ports on different host adapters.
- Use multiple physical adapters.
- Do not use all the ports on each host adapter.

A maximum of two host adapter cards per I/O bay can be used on a DS8800. For more information on adapter cards, see System adapter identification.

Parent topic: Host attachment
Parent topic: Host attachment general considerations

Related concepts
Host attachment general considerations

Related information
Apple Macintosh host attachment
Fujitsu PrimePower host attachment
HP OpenVMS host attachment
HP Tru64 UNIX host attachment
HP-UX host attachment
IBM NAS Gateway 500 host attachment
IBM SAN Volume Controller host attachment
IBM Power Systems host attachment
IBM System z or IBM S/390 host attachment
Linux (x86 or x86_64) host attachment
Microsoft Windows Server host attachment
VMware ESX Server host attachment
Novell NetWare host attachment
Silicon Graphics host attachment
Sun host attachment