

CORAIID
EtherDrive® Mirror Appliance
Configuration Guide

Document Revision 20091112

Table of Contents

1. Introduction.....	4
2. CORAIID EtherDrive SAN Storage.....	5
3. Fault-Tolerant Power Connectivity.....	6
3.1. Power Connections.....	6
3.2. Serial Cable.....	6
4. Network Connectivity.....	6
4.1. Primary Mode.....	6
4.2. Secondary Mode.....	7
4.3. Multi-Path.....	7
5. Management.....	7
5.1. Monitor and Keyboard.....	7
5.2. CEC.....	7
6. Primary Mode.....	8
6.1. Initial configuration.....	8
6.2. Creating a mirror.....	8
6.3. How different sized LUNs are handled.....	10
6.4. When a LUN failure occurs.....	10
6.5. Recovering from a storage array failure.....	10
7. Secondary Mode.....	11
7.1. Initial Configuration.....	11
7.2. Protecting other EMs.....	12
7.3. When a protected EM failure occurs.....	12
7.4. Recovering from a protected EM failure.....	12
8. LUN Masking.....	13
9. Partitions.....	14
9.1. Prepare a mirror LUN for partitioning.....	14
9.2. Creating a partition.....	14
9.3. Label a partition.....	15
9.4. View the partitions on a mirror LUN.....	15
9.5. View the mirror LUN.....	15
10. Firmware Upgrades.....	15
10.1. Connect to EM and create firmware update LUN.....	16
10.2. Mask and online the LUN.....	16
10.3. Copy the firmware to the new EM LUN.....	16
10.4. Reboot the EM and verify the new firmware release.....	17
11. Frequently Asked Questions.....	18
Appendix.....	19
A-1. Primary Mode EM Configuration (without comments).....	19
A-2. Secondary Mode EM Configuration (without comments).....	20
A-3. EM Firmware Upgrade (without comments).....	20

Copyright © 2009 CORAID, Inc.

CORAID, EtherDrive, and RAIDshield are registered trademarks or trademarks of CORAID, Inc. in the United States and/or other jurisdictions.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

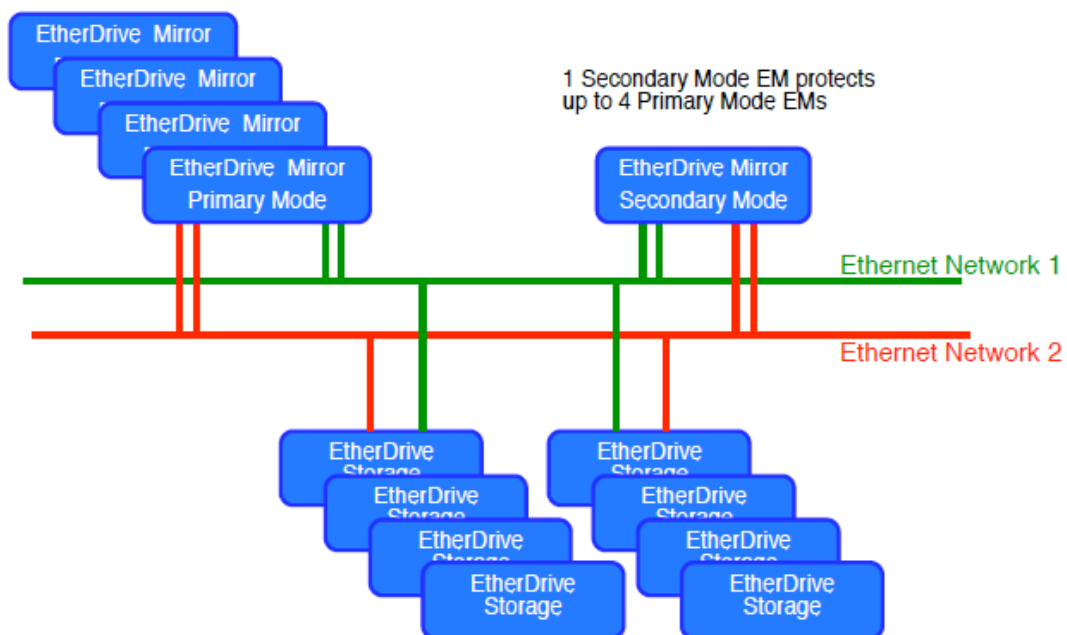
All other marks and names mentioned herein may be trademarks of their respective companies.

1. Introduction

The CORAIID EtherDrive® EM Storage Mirror Appliance is a network appliance for use with the EtherDrive® SAN storage solution. The EM is an “in-band” network appliance that utilizes the open ATA-over-Ethernet (AoE) lightweight storage area network protocol to provide block level LUN mirroring and partitioning. The EM enables LUNs located on an EtherDrive SAN storage appliance to be synchronously mirrored to the same storage chassis or a separate storage chassis. This allows users to build storage systems with No Single Point of Failure (NSPOF). The EM storage mirror takes full advantage of the power and simplicity of the AoE storage protocol to perform fast “on the fly” block address translation of read/write requests sent to EtherDrive® SAN storage.

The EM creates new LUNs by mirroring and partitioning physical LUNs located on one or more EtherDrive® SAN Storage appliances. An EM can create and manage up to 255 separate LUNs. Each of the 255 LUNs can be either a mirrored pair of physical LUNs from EtherDrive storage appliances or a partition created from a mirror LUN. VMware and Linux servers connect to EM LUNs via standard Ethernet connections. EM LUNs look like ordinary physical LUNs to the host server. The EM storage mirror is connected to the SAN allowing it to create mirrored LUNs from any pair of physical LUNs anywhere on the SAN. Members of a mirrored pair can be individual disks (JBOD) or RAID volumes (physical LUNs) located in separate EtherDrive® storage appliances. Once an EM mirror LUN has been created, it can be accessed on the SAN as a single LUN. If smaller LUNs are desired, the mirrored LUN can be partitioned by the EM. Each partition becomes a new LUN available on the SAN.

A fully redundant storage network can be accomplished by using two or more EMs. An EM configured in Secondary Mode can protect up to four Primary Mode EMs. The Secondary EM uses a heartbeat signal sent via the SAN network to monitor all Primary EMs. The Secondary will take over for a Primary when the Primary fails to send a heartbeat. Additionally, during the process of taking over the Primary role, the Secondary EM will remove power from the failed Primary EM – this process eliminates the possibility of a split-brain scenario where two EMs may simultaneously attempt to operate as the same Primary.



Requirements

- Linux Servers: A Linux server must use a software AoE driver in order for it to be able to communicate with a CORAIID EtherDrive appliance. The driver must be version aoe6-71 or newer. The AoE driver may be downloaded from: <http://support.coraid.com/support/linux>
- VMware Servers: A CORAIID EtherDrive HBA must be installed in a VMware server in order for it to be able to communicate with a CORAIID EtherDrive appliance. All CORAIID EtherDrive SAN storage appliances must be upgraded to firmware version **20090303** or higher in order to operate correctly with an EtherDrive HBA. The latest CORAIID EtherDrive SAN appliance firmware release may be downloaded from: <http://support.coraid.com/support/sr/index.html>. The EtherDrive HBA driver may be downloaded from: <http://support.coraid.com/support/esx>
- Since the AoE protocol is not routable, the CORAIID EtherDrive HBA and the CORAIID EtherDrive SAN storage must be connected to the same network broadcast domain.
- The CORAIID EtherDrive EM and EtherDrive SAN must be connected to a network switch that is utilizing jumbo frame support (9000 MTU).

Recommendations

The SAN storage network should be isolated from other network traffic. This allows full utilization of the network for SAN storage.

Maximum Configuration

- Maximum # of LUNs: 255 LUNs
- Maximum # of Mirrored LUNs: 255 LUNs, when no Partition LUNs are created
- Maximum # of Partition LUNs: 254, when only one Mirror LUN is created
- Maximum Mirror LUN Size: Limited only by physical LUN member size

2. CORAIID EtherDrive SAN Storage

Before CORAIID EtherDrive SAN storage can be mirrored using an EM, two LUNs must be configured on one or more EtherDrive SAN storage appliances. The EtherDrive SAN Software Reference Manual should be used as the complete reference for EtherDrive SAN appliance configuration and may be downloaded from:

<http://support.coraid.com/support/sr/index.html>. Below is an example of how a LUN may be setup on an EtherDrive SAN storage appliance. For this example, one LUN is created on storage Shelf 21 and a second LUN is created on storage Shelf 22. Later, the EM will be configured to mirror the LUN on Shelf 21 to the LUN on Shelf 22.

EtherDrive Storage Appliance Shelf 22:

```
SR shelf unset> shelf 22
SR shelf 22> make 5 raid5 22.0-3
SR shelf 22> online 5
```

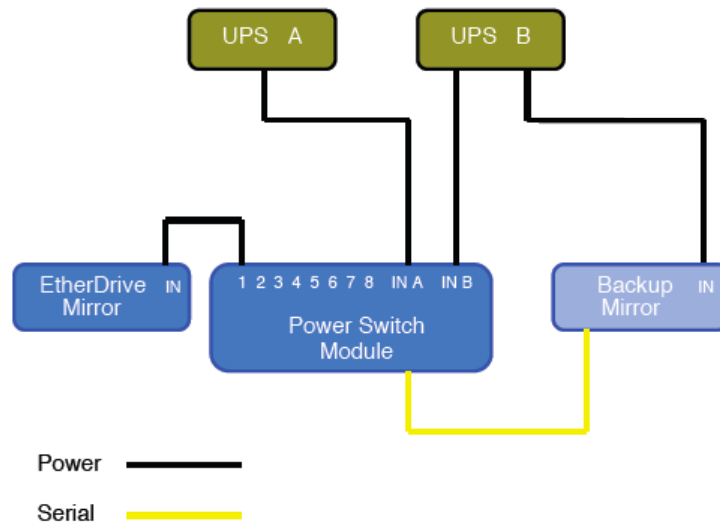
Repeat the same steps for EtherDrive Storage Appliance Shelf 23:

```
SR shelf unset> shelf 23
SR shelf 23> make 5 raid5 23.0-3
SR shelf 23> online 5
```

NOTE: RAID performance is degraded while the RAID on the LUN is being initialized! Full performance can only be achieved once the RAID initialization is complete.

3. Fault-Tolerant Power Connectivity

A fault tolerant configuration can be configured when using two or more EM devices. The EM-FT kit includes an 8 port Power Switch Module. An EM configured for Secondary Mode communicates with the Power Switch Module via a serial connection. When the Secondary Mode EM detects that a protected Primary Mode EM has failed, the Secondary Mode EM instructs the Power Switch Module, via the serial connection, to turn off the configured power port(s) associated with the failed Primary Mode EM.



NOTE: The diagram above depicts the EM-22 power connections. The EM-32 is equipped with two power supplies.

3.1. Power Connections

The Secondary Mode EM and the Power Switch Module should be connected directly to the UPS power source(s). Using two power sources for the Secondary Mode EM and the Power Switch Module will add additional redundancy to the overall storage system.

All Primary Mode EMs should be plugged into one (or two for EM-32) of the power receptacles on the Power Switch Module. Note the outlet number(s) used for each EM as this information will be required in a later step.

3.2. Serial Cable

The serial cable provided with the EM-FT kit should be connected to both the Secondary Mode EM and the Power Switch Module. The serial connection allows the EM to communicate with the Power Switch Module via the Power Switch Module's serial interface. Only Secondary Mode EMs will need to have a serial connection to the Power Switch Module.

4. Network Connectivity

4.1. Primary Mode

Each EM has four Ethernet ports. Each port should be connected to the SAN to ensure the highest level of throughput and redundancy. Additionally, the use of two network switches will add another level of network redundancy.


```
[root@ CORAID~]# cec -s 34 eth1
Probing for shelves ... shelf 34 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM 34>
```

6. Primary Mode

An EM in Primary Mode allows the administrator to create mirrored LUNs and, if desired, partitions of mirrored LUNs. The mirror will consist of two LUNs located either on one EtherDrive Storage Appliance or across two EtherDrive Storage Appliances. Please see Appendix A-1 for the command sequence without additional comments.

6.1. Initial configuration

An EM's default shelf number is -1. CEC is used in the following example to connect to the EM. Alternatively, a monitor and keyboard connected directly to the EM could be used. The first three steps are to connect to the EM, set the mode of the EM to Primary Mode, and then set the shelf address.

6.1.1. Power on and connect to the EM

First, power on the EM and then use CEC to connect to it.

```
[root@ CORAID~]# cec -s -1 eth1
Probing for shelves ... shelf -1 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM -1>
```

6.1.2. Set Primary Mode

After connecting, set the mode to Primary Mode by typing 'setprimary'.

```
EM -1> setprimary
```

6.1.3. Set shelf address

The EM requires a unique shelf address. The valid range of addresses is 1-65000.

```
EM -1> shelf 34
EM 34>
```

6.2. Creating a mirror

The EM will display all of the available EtherDrive LUNs connected to the SAN. The EM can be configured to mirror any EtherDrive LUN to any other EtherDrive LUN. Once the mirror is created, the EM will begin the process of mirroring all data from the source LUN to the mirrored LUN. The initial process of copying the data on the source LUN to the mirrored LUN is run as a background task. This allows immediate use of the newly created mirror LUN.

6.2.1. View available EtherDrive Storage LUNs

The 'Isaoc' command lists all of the available EtherDrive LUNs on the SAN. A LUN cannot be used in the creation of a mirror if it is not present in the TARGET LUN list.

```
EM 34> lsaoe
TARGET          SLENGTH          ADDR          PORTS
22.5            7,814,100,477   0030488a271f  0x0000000c
23.5            7,814,100,477   003048d4012f  0x0000000c
```

6.2.2. Create the mirror

Use the 'mkmir' command to create a mirror. The 'mkmir' command requires arguments that define the new LUN number for the mirror, the source LUN to be mirrored, and the mirrored LUN.

```
EM 34> mkmir 5 22.5 23.5
```

Where:

- 5 Mirror LUN Number: defines the LUN number used to access the mirror. The EM will export the mirrored LUN using the EM's shelf address and this LUN number. In this example, the mirror is exported as 34.5 to the SAN.
- 22.5 Mirror Source LUN Number: defines the EtherDrive LUN that is to be mirrored. When the mirror is initially created, the contents of this LUN will be copied to the mirror destination LUN. The EM will keep all new data synchronized between the mirror source and mirror destination LUNs.
- 23.5 Mirror Destination LUN Number: defines the LUN number that will serve as the mirror destination LUN. When the mirror is initially created, the contents of the mirror source LUN will be copied to this LUN. The EM will keep all new data synchronized between the mirror source and mirror destination LUNs.

6.2.3. Enable the mirror

Once the mirror has been created it will not be accessible to the SAN until it has been enabled. Use the 'online' command to make the mirror LUN visible to the SAN. Then, use the 'lsmir' command to verify the LUN is online. The 'R' next to the target LUN indicates that the contents of the source LUN are being copied, or mirrored, to the target LUN.

```
EM 34> online 5
EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB  SLENGTH  PARTS  LABEL
5      ON      22.5     23.5R    7,814,100,477  n/a
```

Indicates the LUN is being copied to 

6.2.4. Label the mirror

```
EM 34> label 5 Sales1
EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB  SLENGTH  PARTS  LABEL
5      ON      22.5     23.5R    7,814,100,477  n/a  Sales1
```

6.2.5. Check the mirror build status

The status of the initial mirror copy may be viewed by issuing the wstat command.

```
EM 34> wstat
LUN          ACTION  DONE
5            mirror 22.5->23.5  9.28%
```

Once the mirror has completed being built, the 'R' next to the target LUN is removed.

```
EM 34> lsmir
LUN ONLINE   TARGETA   TARGETB   SLENGTH  PARTS  LABEL
5         ON       22.5      23.5     7,814,100,477  n/a  Sales1
```

6.3. How different sized LUNs are handled

Ideally both of the EtherDrive LUNs that comprise the mirror will be the same size and contain the same model hard drives. If the two LUNs that comprise the mirror are not identical in size, the size of the smaller LUN will be used for the mirror size and the remainder of the larger LUN will not be available for other use. For example, if a mirror consists of a 13TB Mirror Source LUN and a 10TB Mirror Destination LUN, the mirror size will be 10TB and the remaining 3TB of the Mirror Source will not be useable.

6.4. When a LUN failure occurs


When either of the EtherDrive LUNs that comprise the mirror fails, the EM will recognize the failure and mark the failed LUN as down. The EM will continue to write/read data to/from the functioning LUN. The mirror will operate in this manner until the failed LUN is corrected, brought back online, and a mirror rebuild is manually initiated.

6.5. Recovering from a storage array failure

When a LUN that comprises part of a mirror fails, the failure should be corrected in a typical manner for the existing failure. For example, if a LUN fails due to one or more disk failures, the failed disks should be repaired or replaced as necessary and then the LUN should be rebuilt. Once the LUN has completed the rebuild process, the mirror should be rebuilt as well. The mirror is rebuilt by issuing the 'rebuild' command from the EM appliance.

In the example below, LUN 23.5 has failed as indicated by the * next to the Shelf.LUN number.

```
**EM 34> lsmir
LUN ONLINE   TARGETA   TARGETB   SLENGTH  PARTS  LABEL
5         ON       22.5      23.5*    7,814,100,477  n/a  Sales1
```

 Indicates this remote LUN has failed

After the cause of the failure has been corrected and the LUN has successfully been brought back online, issuing the rebuild command will begin the re-mirroring process. Two parameters must be supplied to the rebuild command. The first parameter defines the mirror to be rebuilt and the second parameter defines the LUN to be rebuilt. In the example below, the EM mirror LUN 5 will be rebuilt. More specifically, the entire contents of LUN 22.5 will be copied to LUN 23.5.

```
**EM 34> rebuild 5 23.5
2009.10.10 02:45:10 starting rebuild on lun 5: 22.5->23.5
```

After the rebuild command is issued, the * next to the failed LUN is replaced by an 'R' to indicate that the mirror is being rebuilt.

```
EM 34> lsmir
LUN ONLINE TARGETA TARGETB SLENGTH PARTS LABEL
5 ON 22.5 23.5R 7,814,100,477 n/a Sales1
```

 Indicates this LUN is being re-mirrored

The mirror LUN may be used during this time, but there will be a slight degradation in performance capabilities while the mirror is being rebuilt. Once the mirror has completed the rebuild process, the 'R' next to the previously failed LUN will be removed.

7. Secondary Mode

An EM configured for Secondary Mode can protect other EMs configured for Primary Mode. A Secondary Mode EM can be configured to protect up to four other Primary Mode EMs. Once configured, a Secondary Mode EM will assume full responsibility for a failed Primary Mode EM. Please see Appendix A-2 for the command sequence without additional comments.

7.1. Initial Configuration

An EM's default shelf number is -1. CEC is used in the following examples to connect to the EM. Alternatively, a monitor and keyboard connected directly to the EM could be used. The first three steps are to connect to the EM, set the mode of the EM to Secondary Mode, and then set the shelf address

7.1.1. Power on and connect to the EM

First, power on the EM and then use CEC to connect to it.

```
[root@ CORAIID~]# cec -s -1 eth1
Probing for shelves ... shelf -1 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM -1>
```

7.1.2. Set Secondary Mode

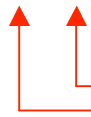
After connecting, set the mode to Secondary Mode by typing 'setsecondary'.

```
EM -1> setsecondary
```

7.1.3. Set shelf address

The EM requires a unique shelf address. The valid range of addresses is 1-65000. When an EM is configured for Secondary Mode and the shelf address is set, two shelf addresses separated by a colon are listed. The first shelf address is the assigned shelf address. The second shelf address is a -1. When a Secondary Mode EM takes over for a Primary Mode EM the second shelf address (normally a -1) will be replaced by the Primary Mode EM's shelf address.

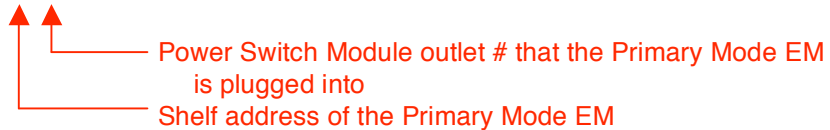
```
EM -1> shelf 35
EM 35 : -1>
```

 Shelf # of failed EM (-1 indicates no failure)
Assigned shelf address

7.2. Protecting other EMs

The Secondary Mode EM must be configured to protect other Primary Mode EMs. Once the Secondary Mode EM has been configured to protect, or watch, a Primary Mode EM, the Secondary Mode EM and the Primary Mode EM exchange heartbeat messages. The outlet number that the Primary Mode EM is plugged into on the Power Switch Module is required for this step. For this example, the Primary Mode EM configured in Section 6 will be protected by adding the following configuration to the Secondary Mode EM:

```
EM 35 :-1> watch 34 1
```



```
EM 35:-1> lswatch
PRIMARY PLUG0 PLUG1 STATUS
34          1          fd=1 iam=0 time=0
```

7.3. When a protected EM failure occurs

Once properly configured, a Secondary Mode EM exchanges heartbeats with all of the configured Primary Mode EMs. When the Secondary Mode EM recognizes that a Primary Mode EM is no longer participating in the heartbeat messages, the Secondary Mode EM will send a command via the serial connection to the Power Switch Module instructing the Power Switch Module to turn off the power port associated with the failed Primary Mode EM. The Secondary Mode EM will then assume all responsibilities of the failed Primary Mode EM.

In the following example, the Primary Mode EM (Shelf 34) has failed and the Secondary Mode EM (Shelf 35) has assumed the role of Primary Mode for the failed EM.

```
** EM 35:34> lsmir
LUN ONLINE TARGETA TARGETB SLENGTH PARTS LABEL
5      ON      22.5    23.5R    7,814,100,477 n/a Sales1
```

7.4. Recovering from a protected EM failure

The first step to recovering from a protected EM failure is to diagnose the cause of the failure. When an EM failure occurs it is most commonly associated with loss of power. In this case, for example, the cause of the power loss needs to be corrected. Once the cause of the EM failure has been corrected, the Secondary Mode EM must be instructed to allow the failed Primary Mode EM to be powered back on. Once the 'unfail' command has been issued on the Secondary Mode EM, the Secondary Mode EM will send a command via the serial connection to the Power Switch Module instructing the Power Switch Module to turn back on the power port associated with the failed Primary Mode EM. The Primary Mode EM will power up and resume control of the mirror LUN. The Secondary Mode EM will resume the duties of watching for another Primary Mode failure.

In the following example, the 'unfail' command is issued on the Secondary Mode EM. After the 'unfail' command is issued, the Primary Mode EM will be powered back on and take back control of the mirror. The Secondary Mode EM will resume the role of monitoring Primary Mode EMs.

```
** EM 35:34> unfail 34
EM 35:-1>
```

8. LUN Masking

LUN masking allows an administrator to restrict the availability of an EtherDrive LUN to one or more hosts. This is accomplished by assigning one or more MAC addresses to an EtherDrive LUN. Once assigned, an EtherDrive LUN will only be visible to those hosts who's MAC address is defined in the mask list.

For example, an EtherDrive Storage LUN may be masked so that only the Primary Mode and Secondary Mode EMs may see the LUN. This would prevent the underlying storage LUNs that comprise the mirror from being presented to an initiator where they may be accidentally accessed directly and risk corrupting the mirror.

Additionally, an EM LUN may be masked so that it's LUNs are visible only to the desired initiators. When masking an EM LUN, it's important to remember to add the MAC addresses of the Secondary Mode EM.

When using an EM with an EtherDrive HBA for VMware ESX, a Linux initiator will be required for maintenance operations such as firmware upgrades and CEC console access. The Linux initiator must have network access to the storage network and may be a physical server or a Virtual Machine running in the ESX environment. In either case, the Linux server's MAC address must be added to the upgrade LUN mask list.

In the example below, the EtherDrive Storage LUNs on shelf 22 and shelf 23 are masked so that they will only be visible to the Primary and Secondary EM. This will ensure that only the EM mirror LUN will be visible to the initiators on the SAN.

```
SR shelf 22> mask 5 +003048dae9f1
SR shelf 22> mask 5 +003048dae9f0
SR shelf 22> mask 5 +003048ba8283
SR shelf 22> mask 5 +003048ba8282
SR shelf 22> mask 5 +003048ba8258
SR shelf 22> mask 5 +003048ba8259
SR shelf 22> mask 5 +003048daebb8
SR shelf 22> mask 5 +003048daebb9
SR shelf 22> mask
5 003048dae9f1 003048dae9f0 003048ba8283 003048ba8282 003048ba8258
003048ba8259 003048daebb8 003048daebb9
```

```
SR shelf 23> mask 5 +003048dae9f1
SR shelf 23> mask 5 +003048dae9f0
SR shelf 23> mask 5 +003048ba8283
SR shelf 23> mask 5 +003048ba8282
SR shelf 23> mask 5 +003048ba8258
SR shelf 23> mask 5 +003048ba8259
SR shelf 23> mask 5 +003048daebb8
SR shelf 23> mask 5 +003048daebb9
```

```
SR shelf 23> mask
5 003048dae9f1 003048dae9f0 003048ba8283 003048ba8282 003048ba8258
003048ba8259 003048daebb8 003048daebb9
```

9. Partitions

Once a Mirror LUN is created, it may additionally be partitioned into smaller LUNs. Each new partition will be presented by the EM to the storage network as a unique LUN. Partitioning a large Mirror LUN allows larger LUNs to be presented to the SAN as multiple smaller LUNs.

In the example below, two partitions are created from the mirror LUN that was created in section 6.2.2 (LUN 5):

9.1. Prepare a mirror LUN for partitioning

By default, partitions are not enabled for newly created mirror LUNs. Partitioning is enabled on a per mirror LUN basis. Enabling partitions on one mirror LUN does not affect any other mirror LUN. To add a partition to a mirror LUN, the mirror LUN must be in an 'offline' state and partitioning must be enabled on the mirror LUN by issuing the 'allowpart' command. The mirror LUN must remain in an 'offline' state; thus, the original mirror LUN will no longer be accessible to the SAN – only the new partition LUNs will be accessible to the SAN. Once partitioning is enabled on a mirror LUN, the first 1024 blocks of the mirror LUN are reserved to allow space to store partitioning configuration. Enabling partitions on a mirror LUN will cause the loss of all current data residing on the mirror LUN.

```
EM 34> offline 5
EM 34> allowpart 5
Allowing partitions will irrevocably modify the lun.
Are you sure you want to make lun 5 partitionable? [n]: y
```

```
EM 34> lsmir
LUN ONLINE TARGETA TARGETB SLENGTH PARTS LABEL
5 OFF 22.5 23.5 7,814,100,477 0 Sales1
```

9.2. Creating a partition

Partitions are defined based on block addresses – where each block contains 512 bytes. In the example below, the 'mkpart' command is used to create a new partition from mirror LUN 5. The new partition is assigned the LUN # 6 and is defined to begin at the first available block - which is block 1024. The new partition is defined to be 100GB (100,000,000,000/512 = 195312500 blocks).

```
EM 34> mkpart 6 5 1024 195312500
```

Use the 'lspart' command to display the next available block for creating a new partition.

```
EM 34> lspart 5
LUN SBASE SLENGTH LABEL
0 0 1,024 <reserved>
6 1,024 195,312,500
195,313,524 7,618,786,953 <free>
7,618,786,953 Total Free
```



Next available block for creating another partition.

In the example below, the 'mkpart' command is used to create a second partition from mirror LUN 5. The new partition is assigned the LUN # 7 and is defined to begin at the first available block - which is block 195313524. The new partition is defined to be 200GB (200,000,000,000/512 = 390625000 blocks).

```
EM 34> mkpart 7 5 195313524 390625000
EM 34> lspart 5
```

LUN	SBASE	SLENGTH	LABEL
	0	1,024	<reserved>
6	1,024	195,312,500	
7	195,313,524	390,625,000	
	585,938,524	7,228,161,953	<free>
		7,228,161,953	Total Free

 Next available block for creating another partition.

9.3. Label a partition

Using the label command, add labels to the new LUNs to describe their purpose.

```
EM 34> label 6 Sales_database
EM 34> label 7 Sales_tools
```

9.4. View the partitions on a mirror LUN

Use the lspart command to show the details of the partitions on mirror LUN 5.

```
EM 34> lspart 5
```

LUN	SBASE	SLENGTH	LABEL
	0	1,024	<reserved>
6	1,024	195,312,500	Sales_database
7	195,313,524	390,625,000	Sales_tools
	585,938,524	7,228,161,953	<free>
		7,228,161,953	Total Free

9.5. View the mirror LUN

Use the lsmir command to show the details of mirror LUN 5. Notice the mirror LUN is in an offline state and that it now contains 2 'PARTS' – which are the two partitions that were just added.

```
EM 34> lsmir
```

LUN	ONLINE	TARGETA	TARGETB	SLENGTH	PARTS	LABEL
5	OFF	22.5	23.5	7,814,100,477	2	Sales1

10. Firmware Upgrades

EM firmware updates are available from the CORAIID support website (<http://support.coraid.com/support/em/index.html>). The first step to updating the EM firmware is to create a new LUN using the mkflash command on the EM. Once created, the LUN may optionally be masked and then must be onlined. Once the LUN is onlined, the new EM firmware may be copied to the update LUN. Once the firmware has successfully been copied to the new LUN on the EM, the EM must be rebooted for the new firmware to load. Please see Appendix A-3 for the command sequence without additional comments.

In the example below, a Linux initiator is used CEC to the EM and also to copy the new EM firmware image to the EM appliance.

10.1. Connect to EM and create firmware update LUN

First, a CEC connection is established from the Linux server to the EM.

```
[root@ CORAIID~]# cec -s 34 eth1
Probing for shelves ... shelf 34 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM34> release
Running Release: Tue Jul 29 14:45:40 EDT 2008
```

Next, the EM is instructed to boot from the flash partition A.

```
EM34> setboot
will boot from A
```

Finally, a flash LUN is created using the mkflash command. When issuing the mkflash command, a currently unused LUN number must be supplied to the command (in this example, the LUN is assigned the number 15). The mkflash command is also instructed to use boot partition A.

```
EM34> mkflash 15 A
```

10.2. Mask and online the LUN

Masking the LUN used for firmware updates is optional. Masking the LUN will prevent the firmware update LUN from being presented to all of the initiators on the SAN. The use of masking is especially beneficial for a VMware ESX environment where new LUNs are reported to the ESX initiator. The LUN created for firmware updates will not survive reboot of the EM, but ESX will continue to show the LUN as an available target. Masking the LUN with the Linux server's MAC address will prevent this from occurring.

```
EM34> mask 15 +003022d5e671
```

MAC address of the Linux server where the update
is initiated from.

Next, online the LUN using the 'online' command.

```
EM34> online 15
```

```
EM 34> lsun 15
LUN ONLINE LENGTH TYPE LABEL
15 ON 0.000GB F Flash A
```

10.3. Copy the firmware to the new EM LUN

All of the steps in this section are completed on the Linux server. The Linux server used in this example has the AoE driver and Tools installed. For more information about the AoE driver for Linux and AoE Tools, please see <http://support.coraid.com/support/linux/>.

First, ensure that the new LUN is visible to the Linux server.

```
[root@ CORAID~]# aoe-stat | grep 34
e34.15 0.008GB eth0 1024 up
```

Next, verify the md5sum with that reported on the EM support page.

```
[root@ CORAID~]# md5sum EM20080731
945be5f330fdb7944c10673842a8bc5b EM20080731
```

Finally, use the coraid-update command to copy the new firmware to the newly created EM LUN.

```
[root@ CORAID~]# coraid-update EM20080731 /dev/etherd/e34.15
```

Path to EM device _____ ↑
 EM shelf address _____ ↑
 EM LUN # _____ ↑

As an alternative to the coraid-update command, dd may be used to copy the file to new EM LUN. Note, this step is not necessary if the coraid-update command has already successfully been issued.

```
[root@ CORAID~]# dd if=EM20080731 of=/dev/etherd/e34.15
291+1 records in
291+1 records out
149060 bytes (149 kB) copied, 0.219686 s, 679 kB/s
```

10.4. Reboot the EM and verify the new firmware release

The final step in updating the EM firmware is to reboot the EM by issuing the 'reboot' command to the EM. The release information is provided as the EM boots. Verify the release information matches the version that was just installed.

```
EM34> reboot
rebooting...
Coraid Boot ROM
Boot A?
booting...
Done
2,145,779,712 bytes of memory
EM: CORAID EtherDrive Storage Mirror Appliance
EM Release: Thu Jul 31 12:38:28 EDT 2008 ← Release information
BIOS build 11/02/07
2,145,779,712 bytes
Flash size 32 MB
Console password unset. Access granted.
EM34>
```

11. Frequently Asked Questions

1. Is the Secondary Mode EM a separate product?

No. Any EM can be configured to operate in Primary Mode or Secondary Mode. At least two EMs are required for a fault tolerant setup. At least one EM will operate in Primary Mode and one EM will operate in Secondary Mode.

2. Can I just use a Primary Mode EM or do I have to additionally use a Secondary Mode EM?

The use of an EM in Secondary Mode is optional. An EM in Secondary Mode provides fault tolerance for up to 4 EMs operated in Primary Mode.

3. What do I have to configure to use Multi-path?

Nothing! All EtherDrive products, including the EM, automatically handles Multi-path for all of the available ethernet ports.

4. Can all model EtherDrive storage arrays be used with an EM?

All of the latest EtherDrive models are compatible with the EM. However, some of the earlier models are not recommended for use with the EM. The SR1520 and SR420 are not recommended for use with the EM due to the fact that only one of the two Ethernet ports supports a MTU of 9000. Additionally, the SR421 is not compatible with the EM due to none of the available Ethernet ports supporting a MTU of 9000.

5. I have added a new mirror LUN to my EtherDrive EM appliance, how do I get it to show up in the vSphere Client?

You must perform a storage rescan from vSphere Client in order for newly available LUNs to be presented. This is accomplished by clicking 'Rescan' from the 'Storage Adapters' page under the 'Configuration' tab.

Appendix

A-1. Primary Mode EM Configuration (without comments)

```
[root@ CORAIID~]# cec -s -1 eth1
Probing for shelves ... shelf -1 found.
connecting ... done.
Escape is Ctrl-\

EM -1>
EM -1> setprimary
EM -1> shelf 34
EM 34>
EM 34> lsaoc
TARGET      SLENGTH      ADDR          PORTS
22.5        7,814,100,477 0030488a271f 0x0000000c
23.5        7,814,100,477 003048d4012f 0x0000000c

EM 34> mkmir 5 22.5 23.5
EM 34> online 5
EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB      SLENGTH  PARTS  LABEL
5      ON      22.5      23.5R      7,814,100,477  n/a
EM 34> label 5 Sales1
EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB      SLENGTH  PARTS  LABEL
5      ON      22.5      23.5R      7,814,100,477  n/a  Sales1
EM 34> wstat
LUN          ACTION  DONE
5            mirror 22.5->23.5 9.28%
EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB      SLENGTH  PARTS  LABEL
5      ON      22.5      23.5      7,814,100,477  n/a  Sales1

EM 34> offline 5
EM 34> allowpart 5
Allowing partitions will irrevocably modify the lun.
Are you sure you want to make lun 5 partitionable? [n]: y

EM 34> lsmir
LUN ONLINE  TARGETA  TARGETB      SLENGTH  PARTS  LABEL
5      OFF     22.5      23.5      7,814,100,477  0  Sales1

EM 34> mkpart 6 5 1024 195312500
EM 34> lspart 5
LUN          SBASE          SLENGTH LABEL
6            0              1,024 <reserved>
            1,024          195,312,500
            195,313,524          7,618,786,953 <free>
                                7,618,786,953 Total Free

EM 34> mkpart 7 5 195313524 390625000
EM 34> lspart 5
```

LUN	SBASE	SLENGTH	LABEL
	0	1,024	<reserved>
6	1,024	195,312,500	
7	195,313,524	390,625,000	
	585,938,524	7,228,161,953	<free>
		7,228,161,953	Total Free

EM 34> label 6 Sales_database

EM 34> label 7 Sales_tools

EM 34> lspart 5

LUN	SBASE	SLENGTH	LABEL
	0	1,024	<reserved>
6	1,024	195,312,500	Sales_database
7	195,313,524	390,625,000	Sales_tools
	585,938,524	7,228,161,953	<free>
		7,228,161,953	Total Free

EM 34> lsmir

LUN	ONLINE	TARGETA	TARGETB	SLENGTH	PARTS	LABEL
5	OFF	22.5	23.5	7,814,100,477	2	Sales1

A-2. Secondary Mode EM Configuration (without comments)

```
[root@ CORAID~]# cec -s -1 eth1
Probing for shelves ... shelf -1 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM -1>
EM -1> setsecondary
EM -1> shelf 35
EM 35 : -1>
EM 35 : -1> watch 34 1
```

A-3. EM Firmware Upgrade (without comments)

```
[root@ CORAID~]# cec -s 34 eth1
Probing for shelves ... shelf 34 found.
connecting ... done.
Escape is Ctrl-\
```

```
EM34> release
Running Release: Tue Jul 29 14:45:40 EDT 2008
EM34> setboot
will boot from A
EM34> mkflash 15 A
EM34> mask 15 +003022d5e671
EM34> online 15
EM 34> ls lun 15
LUN ONLINE LENGTH TYPE LABEL
15 ON 0.000GB F Flash A
[root@ CORAID~]# aoe-stat | grep 34
```

```
e34.15 0.008GB eth0 1024 up  
[root@ CORAID~]# md5sum EM20080731  
945be5f330fdb7944c10673842a8bc5b EM20080731  
[root@ CORAID~]# coraid-update EM20080731 /dev/etherd/e34.15
```

```
EM34> reboot  
rebooting...  
Coraid Boot ROM  
Boot A?  
booting...  
Done  
2,145,779,712 bytes of memory  
EM: CORAID EtherDrive Storage Mirror Appliance  
EM Release: Thu Jul 31 12:38:28 EDT 2008  
BIOS build 11/02/07  
2,145,779,712 bytes  
Flash size 32 MB  
Console password unset. Access granted.  
EM34>
```