

Dell Fault Resilient Memory

A Dell Technical White Paper

December 2013

Avinash Bendigeri

Shiva Katta

Krishnaprasad K

Dell Hypervisor Engineering



THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo are trademarks of Dell Inc. VMware, ESXi are trademarks of VMware, Inc. in the United States and/or other jurisdictions. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

December 2013

Table of Contents

Introduction to Dell Fault Resilient Memory (FRM).....	4
Benefits of enabling FRM.....	4
Pre-requisites.....	4
Audience and Scope.....	5
Enabling FRM in PowerEdge servers.....	5
Enabling and verifying ReM from VMware ESXi	6
Monitoring FRM Redundancy Failure.....	7
Summary	8
References.....	8

Table of Figures

Figure 1 : Enabling FRM Memory Operating Mode in the BIOS	6
Figure 2: FRM Token in ESXi.....	6
Figure 3: ESXi esxcli Command Reference	7
Figure 4: Listing ESXi ReM kernel parameter.....	7
Figure 5: FRM Memory redundancy lost when uce is injected to memory channels in Socket0.....	7
Figure 6: Multi-bit memory error SEL entries when uce is injected to memory channels in Socket1	8

Introduction to Dell Fault Resilient Memory (FRM)

Fault Resilient Memory (FRM) is a new Memory Operating Mode available on Dell PowerEdge 12th generation servers. The mode establishes an area of memory that is fault resilient and can be used by an Operating System supporting the FRM feature to load critical Applications/Services & the OS Kernel to maximize system availability. VMware ESXi 5.5 supports FRM through its Reliable Memory (ReM) Feature.

Benefits of enabling FRM

FRM creates a highly resilient memory zone for the hypervisor, protecting it from all but the most severe of memory errors. With the VMware Reliable Memory feature, vSphere v5.5 is able to take advantage of this zone, providing the hypervisor with a strong protection from memory faults that would otherwise bring down the entire system. When used in conjunction with other Dell PowerEdge reliability and redundancy features for virtualization -- such as the Internal Dual SD Modules usable with ESXi installations and Memory Page Retire, which helps the hypervisor to locate and fence off areas prone to many correctable errors that could lead to an uncorrectable error in the standard memory space -- it creates a highly reliable VMware virtualization server environment for customers concerned with high availability and uptime.

Pre-requisites

The pre-requisites for enabling FRM are as below

1. Dell PowerEdge 12th generation server(s) supporting FRM
 - ✓ R720
 - ✓ R720-XD
 - ✓ R620
 - ✓ M620
 - ✓ T620
2. A unique Dell BIOS revision is required to support FRM feature. The BIOS version should be 2.0.19 or later.
3. DIMMs populated to support Memory Mirror Mode. For further details, refer to respective Server Hardware Owner's Manual.
4. VMware vSphere v5.5 or later offering the Reliable Memory feature (Enterprise or Enterprise Plus editions).

Audience and Scope

The intended audience for this whitepaper includes IT Administrators and Channel Partners planning to make use of this new memory operating mode introduced in the 12th generation Dell PowerEdge servers. Fault Resilient Memory is a feature aimed at customers concerned with highly reliable virtualization platforms. It is not targeted to customers leveraging highly overcommitted memory resources, or needing the highest possible memory performance.

Enabling FRM in PowerEdge servers

On a Supported 12th generation Dell PowerEdge server, navigate to “Memory Settings” in the BIOS setup and select “Dell Fault Resilient Mode” option under “Memory Operating Mode” section. Save the settings and Exit from the BIOS Setup. Installing VMware ESXi 5.5 on the server makes vmkernel and other critical Applications/Services operate in the reliable region by default. If using an edition of VMware vSphere with the Reliable Memory feature enabled, it will automatically locate the FRM space and load the hypervisor into the protected zone. The user does not have to make any configuration changes within ESXi or vCenter to enable ReM.

NOTE: It is strongly recommended that you do NOT enable FRM on earlier vSphere versions than v5.5, or any other OS. While the system will function, the memory used to create the fault tolerant zone will be still consumed and the protected zone provided will be used randomly, thus causing wastage of system memory.

The below screenshot(s) explain enabling FRM token in server BIOS.

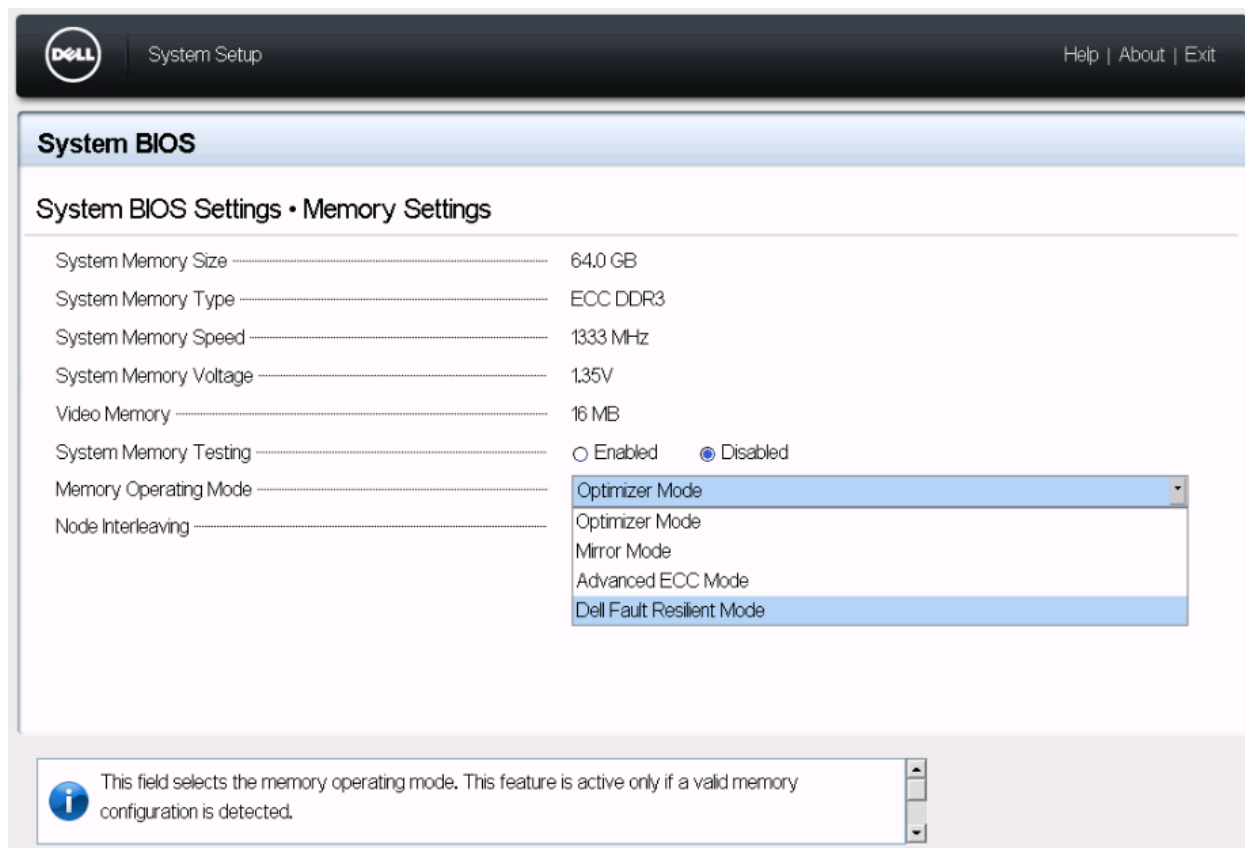


Figure 1 : Enabling FRM Memory Operating Mode in the BIOS

Enabling and verifying ReM from VMware ESXi

VMware ESXi 5.5 enables Reliable Memory (ReM) by default (for supported versions - Enterprise, Enterprise Plus). The user need not make any configuration changes in ESXi to enable ReM. There are few utilities VMware provides to verify the reliable memory region status. For example the command (**esxcli sched reliablemem get**) mentioned in the screenshot indicates that the system is FRM/ReM enabled or not. A value of 'true' indicates that it's enabled.

```
~ #  
~ # esxcli sched reliablemem get  
true  
~ #
```

Figure 2: FRM Token in ESXi

The example screenshot is extracted from a PowerEdge M620 server with both the processors populated and physical memory populated to 128GB.

The command **esxcli hardware memory get** produces the output as shown in the figure 4. This describes the total amount of physical memory available in the system with the amount of reliable memory exposed to the hypervisor as well. The NUMA node count ‘2’ indicates that it’s a dual processor socket system.

```
~ #
~ # esxcli hardware memory get
Physical Memory: 103031480320 Bytes
Reliable Memory: 34310393856 Bytes
NUMA Node Count: 2
~ #
```

Figure 3: ESXi esxcli Command Reference

The command (**esxcli system settings kernel list | grep useReliableMem**) mentioned in the screenshot indicates whether the kernel parameter **useReliableMem** is enabled or not. By default, this parameter is set to TRUE.

```
~ # esxcli system settings kernel list | grep useReliableMem
useReliableMem          Bool      System is aware of reliable memory.
                                TRUE      TRUE      TRUE
~ #
```

Figure 4: Listing ESXi ReM kernel parameter

Monitoring FRM Redundancy Failure

This section details about monitoring the FRM related events from iDRAC System Event Log (SEL). The entry noted in Figure6 pops up in SEL when an uncorrectable soft error occurs to a memory address associated to Socket0. This clearly shows that the FRM provided memory redundancy is lost and occurrence of one more uncorrectable error to any of the memory addresses’ in Socket0 will result in a Purple Screen Of Death (PSOD).

System Event Log

Severity	Date/Time	Description
Instructions: The System Event Log contains information about the managed system. To sort the log by column, click a column header.		
	Thu Nov 21 2013 10:04:27	Memory mirror redundancy is lost. Check memory device at location(s) DIMM_A1.

Figure 5: FRM Memory redundancy lost when uce is injected to memory channels in Socket0

The below entries pops up in SEL when an uncorrectable soft error occurs to a memory address associated to Socket1 even when FRM is enabled (Since FRM provides redundancy in memory channels associated to Socket0 only). Thus resulting for a hypervisor crash.

System Event Log		
Severity ▾	Date/Time ▾	Description ▾
Instructions: The System Event Log contains information about the managed system. To sort the log by column, click a column header.		
✖	Thu Nov 21 2013 10:07:24	Multi-bit memory errors detected on a memory device at location(s) DIMM1,DIMM2,DIMM3,DIMM4,DIMM5,DIMM6,DIMM7,DIMM8.
✖	Thu Nov 21 2013 10:07:20	Multi-bit memory errors detected on a memory device at location(s) DIMM_B1.

Figure 6: Multi-bit memory error SEL entries when uce is injected to memory channels in Socket1

Summary

This paper details both Dell & VMware technologies to make memory reliable for virtual platform. It explains about the feature, the pre-requisites required to enable the reliable memory and the steps needed to enable FRM from the PowerEdge server BIOS. The paper also provides an insight into the benefits that FRM can bring in for the end users.

References

- [What's new in VMware vSphere 5.5 platform](#)
- [FRM : A continuum of innovation in partnership with VMware](#)