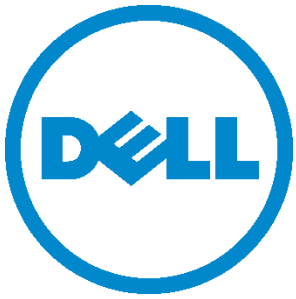

Dell's vFlash SD Flash Media Card – Better Endurance & Performance via Over-Provisioning

Santosh Kumar

November 2011

Table of Contents

Introduction

Over-provisioning (sometimes spelled as OP, over provisioning, or overprovisioning) is the difference between the physical capacity of the Flash memory and the logical capacity presented through the operating system (OS) as available for the user. It is a term used to describe the process of increasing the spare area on a flash memory card/drive. In short, over-provisioning means that more physical NAND flash is present than what the system can access at any time. This document provides details of the Dell's over-provisioned MLC NAND based SD Card.

Over-provisioning increases the endurance and performance of the flash memory card/drive and decreases the flash memory wear and helps lower the write amplification when the flash memory controller writes to the flash memory.

To mitigate this erase-before-write effect, over provisioning is used by Dell to provide pre-erased blocks. Over provisioning allows for the direct writing of data into the over provisioned or "hidden" blocks of space in the foreground operations. In the background, a cleanup routine of moving the data from the "hidden" area into the user area occurs. This cleanup process of freeing up and erasing blocks is managed to ensure that most writes to the drive do not require the slower erase step prior to writing to the sector. Having more "hidden" capacity in "free" blocks available to the SD card allows for better write performance.

Dell's 8GB and 16GB vFlash SD Card

As of Dec 2011, Dell released high endurance 8GB (Giga Byte) and 16GB vFlash SD cards based on MLC NAND flash. These cards are 100% factory overprovisioned for highest endurance.

The 8 GB vFlash SD card is built with two 64 Gbit NAND flash chips totaling physical NAND capacity of 16 GB. It has 8GB more physical NAND flash than the logical capacity presented through the operating system (OS) as available for the user. The extra 8 GB NAND blocks are reserved for bad block replacement in the case of NAND block failure.

The 16 GB vFlash SD card is built with two 128 Gbit NAND flash chips totaling physical NAND capacity of 32 GB. It has 16GB more physical NAND flash than the logical capacity presented through the operating system (OS) as available for the user. The extra 16 GB NAND blocks are reserved for bad block replacement in the case of NAND block failure.

Implementation

NAND flash memory is a consumable device and susceptible to wear, due to repeated program and erase cycles that are commonly done in data storage applications and systems using Flash Translation Layer (FTL). Constantly programming and erasing to the same memory location eventually wears that portion of memory out and makes it invalid. As a result, the NAND flash will have limited lifetime. To prevent scenarios such as these from occurring, special algorithms are deployed within the SSD called wear leveling. As the term suggests, wear leveling provides a method for distributing program and erase cycles uniformly throughout all of the memory blocks within the SSD. This prevents continuous program

and erase cycles to the same memory block, resulting in extending the life to the overall NAND flash memory.

Latest NAND are manufactured by 2x nm and 1y nm technology process node. These NAND devices have very limited program and erase cycles. Especially the consumer grade three level cell (TLC) NAND device, which has program and erase cycles as low as 500 cycles. Even with advance wear leveling the endurance of the drive is very limited.

The overprovisioned NAND blocks are pools of NAND blocks set aside (reserved) for bad block replacement. The figure below illustrates the active block pool, reserved block pool and bad block pool. The active block pool is set of NAND blocks actively used to store the user data. The reserved block pool is set of good fresh NAND blocks reserved for the bad block replacement. The bad block pool is set of NAND block which has failed and gone bad. Upon a block failure in active pool, the failed block is replaced by a good fresh block from the reserved block pool. A failed block is moved to bad block pool adding bad block pool size by one. A fresh good block from the reserved block pool is moved to the active block pool thus maintaining the active block number in active block pool. The number of reserved blocks in the reserved block pool is reduced by one. This process is continued on any block failure in the active block until the fresh good block has been exhausted.

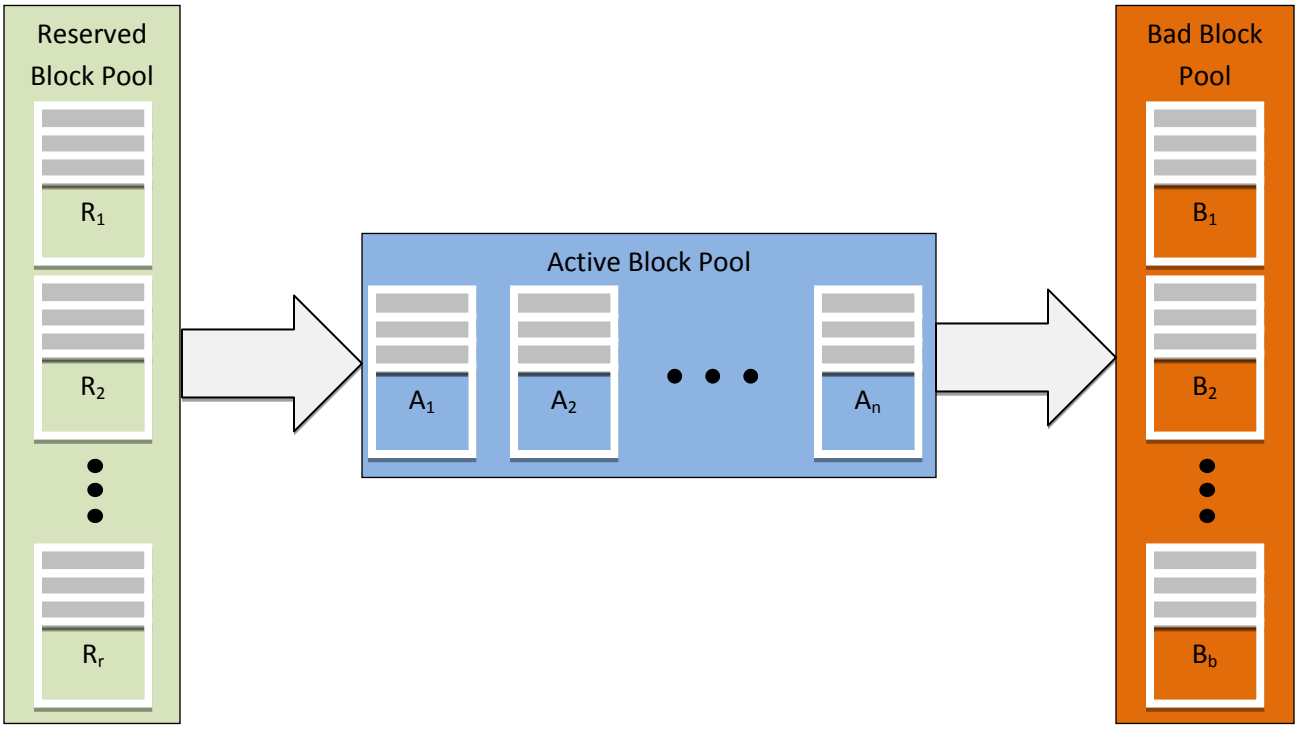

Figure1: Active, reserved, and bad block pool illustration

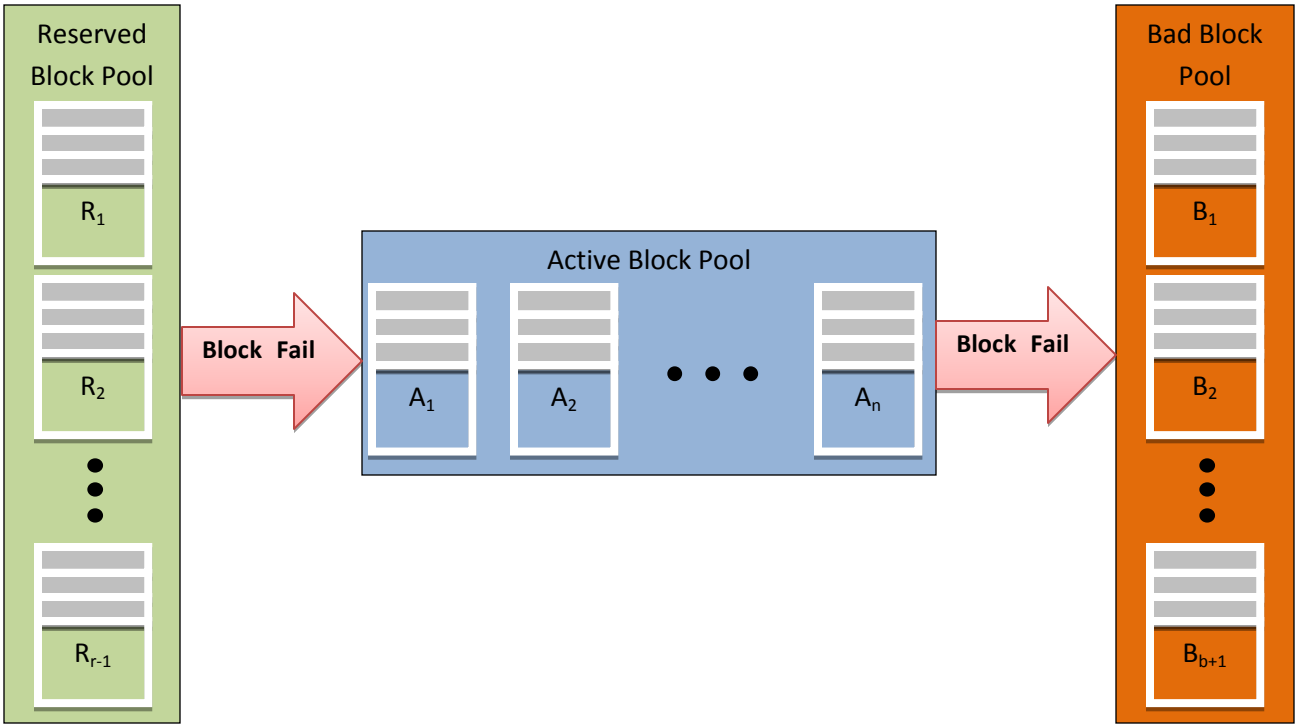

Figure2: Active, reserved, and bad block pool illustration after a block failure

Conclusion

Dell's MLC based 8 GB and 16 GB class 10 SD card has higher endurance and is a higher performance flash media card, designed for enterprise application when compared to consumer/retail grade SD cards.

New for Dell's 12th Generation Servers, customers who purchase iDRAC7 Enterprise have the option to use a locally acquired SD card or they may purchase one from Dell. While SD cards appear the same, not all Class 10 SD cards are the same. The SD cards used by Dell have been enhanced to provide added endurance and performance though over-provisioning.