IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION

RADIAN MEMORY SYSTEMS LLC, *Plaintiff*, v. SAMSUNG ELECTRONICS CO., LTD., AND SAMSUNG ELECTRONICS AMERICA, INC., *Defendants*.	Civil Action No.: 2:24-cv-1073 **JURY TRIAL DEMANDED**

PLAINTIFF'S COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Radian Memory Systems LLC ("Radian") files this complaint under 35 U.S.C. § 271 against Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc., (collectively, "Samsung" or "Defendants") for infringement of U.S. Patent Nos. 11,544,183 ("the '183 Patent"); 11,709,772 ("the '772 Patent"); 11,681,614 ("the '614 Patent"); 11,740,801 ("the '801 Patent"); 11,347,657 ("the '657 Patent"); 11,347,656 ("the '656 Patent"); and 11,307,995 ("the '995 Patent") (each, an "Asserted Patent," and collectively, the "Asserted Patents"). Radian, on personal knowledge as to its own acts, and upon information and belief as to all others based on investigation, alleges as follows:

THE PARTIES

A. Radian

1. Radian Memory Systems LLC ("Radian") is a Texas limited liability company having its principal place of business at 6275 West Plano Parkway, Suite 500, Plano, Texas 75093. Radian is the owner and assignee of patents previously owned by Radian Memory Systems Inc.

("Radian Inc.").

B. Samsung

2. Defendant Samsung Electronics Co., Ltd. ("SEC") is a foreign corporation organized and existing under the laws of the Republic of Korea. SEC has a principal place of business located at 129 Samsung-ro, Yeongtong-Gu, Suwon-Si, Gyeonggi-Do, 443-742, Republic of Korea.

3. Defendant Samsung Electronics America, Inc. ("SEA") is a wholly owned subsidiary corporation of Samsung Electronics Co., Ltd. and is organized and existing under the laws of New York with a principal place of business at 85 Challenger Road, Ridgefield Park, New Jersey 07660. Samsung Electronics America, Inc. has offices and other facilities in Texas, including within this District, at least at 6625 Excellence Way, Plano, Texas 75023; 12100 Samsung Boulevard, Austin, Texas 78754; 1301 E. Lookout Drive, Richardson, Texas 75082, 1000 Klein Road, Plano, Texas 75075; 2601 Preston Road, Suite #1214, Frisco, Texas 75034; 6625 Declaration Way, Plano, Texas 75023; 6105 Tennyson Parkway, Plano, Texas 75203; 1005 Placid Avenue, Suite #120, Plano, Texas 75074; and 6555 Excellence Way, Suite #00100, Plano, Texas 75023.

4. Samsung Electronics America, Inc. may be served through its registered agent for service of process, CT Corporation System, 1999 Bryan St., Suite. 900, Dallas, Texas 75201.

FACTUAL BACKGROUND AND ALLEGATIONS

5. Radian presents its claims and allegations herein in conjunction and/or in the alternative. To the extent that certain claims or allegations may be inconsistent with one another, such allegations should be understood to be made in the alternative.

6. In the interest of providing detailed averments of infringement, Radian has

identified below at least one claim of each of the Asserted Patents to demonstrate infringement. However, the selection of claims should not be considered limiting, and additional claims of the Asserted Patents that are infringed by Samsung will be disclosed in compliance with the Court's schedule.

A. Overview

7. Radian Memory Systems pioneered a new paradigm for managing flash storage media that it called Cooperative Flash Management ("CFM"). The company began research and development efforts for the technology in 2010 to overcome the inefficiencies of flash translation layers (or "FTLs")—the foundation of most flash memory solid-state drive (or flash "SSD") ever produced before—while still providing the ease of integration and data center reliability of FTL-based implementations that had not otherwise been possible. Radian's innovations have been adopted and deployed in the industry, including by Samsung.

8. At the time of the inventions, next generation data centers needed storage solutions that were deterministic, had low latency, and were cost-effective. Flash SSDs were a key building block in realizing this type of storage solution. Flash SSDs at the time were based on a FTL that provided an abstraction to manage the flash memory and communicate with host systems (e.g., an operating system on a server). However, the abstraction provided by the FTLs was redundant and inefficient, so much so that it led to expensive overhead, unpredictable latency spikes, suboptimal performance, and premature wear-out of the flash media. These drawbacks undermined the potential for superior performance and cost-savings that could have come from flash SSD solutions at the time.

9. NAND flash memory is a non-volatile media that has specific characteristics and attributes. Using NAND as part of a storage solution means that those attributes and programming

operations specific to NAND have to be managed. This management was typically performed by the SSD at the cost of the inefficiencies described above. At the same time, a host had components like file systems, block virtualization managers, or object/key value stores, that provided the capabilities to manage storage media. But if the host tried to either program data into or read data directly from raw NAND flash, it would need significant customization to its own system software based on the specifics of the SSD hardware that it was attempting to access. That would make the host dependent on a particular set of SSD hardware or a particular SSD configuration. The host would also have to expend its own resources to take on lower-level NAND maintenance and management operations that would otherwise be handled by the SSD.

10. Users, such as storage system array suppliers or operators of data centers, were faced with a dilemma. They could either invest significant time and money into customizing their host software and SSD hardware to avoid redundancies and inefficiencies from FTLs. But doing so would then leave them exposed to the SSD hardware becoming obsolete in the future, which would necessitate a new round of significant resource expenditure to customizing their system software for the next generation of hardware. The other option was to adapt the host system's usage of the SSD storage systems to the "one size fits all" regime dictated by off-the-shelf SSDs. But as described above, that would cause the overall system to incur significant costs in terms of inefficiencies and suboptimal performance.

11. Against this backdrop, Radian introduced CFM, a type of software-defined flash (or "SDF"), and its Symphonic-branded CFM implementation. Symphonic abstracted some of the lower-level flash attributes for the host and provided for a redistribution of host and SSD responsibilities. This enabled the SSD to avoid the inefficiencies of the conventional FTL while providing improved quality-of-service, performance, cost, endurance, and data center class

product functionalities.

12. Radian's technological innovations provided step-change improvements in SSD solutions. Radian began to introduce and promote its technology to the industry, including sponsoring trade shows, hosting exhibit booths, providing technology presentations, and demonstrations at industry conferences. Radian's innovations were recognized and praised in the industry. Major industry publications such as EE Times, The Register, Tom's Hardware, and CRN covered these efforts. Radian also won the award for "Most Innovative Technology" at the industry's largest conference.

13. Eventually, the industry converged on Radian's solutions. This included industry-wide non-volatile memory standard specifications, such as the NVM Express Zoned Namespaces Command Set Specification, that was published with interfaces and commands based on Radian's technology and inventions. Despite this, however, Radian itself was cast out by the industry. Perhaps this was retaliation for Radian's refusal to join the NVMe industry standard organization and associated Technical Working Groups, which would have required Radian to give the industry royalty-free licenses to any of its patents that were practiced by implementations of the standards specifications. Samsung was a key participant in the industry, serving not only as a member of the NVMe industry standard organization, but also as a member of the specific Technical Working Group for the Zoned Namespaces Specification.

14. The inventions within Radian's various innovations are protected by Radian's patent portfolio, including the Asserted Patents in this case. Samsung infringes the Asserted Patents as more fully described below.

B. Background on Conventional Solid-State Drives ("SSDs")

15. This section provides background on conventional SSDs. While some of these

features appear in the Accused Instrumentalities, not all of them do. In fact, Radian's inventions have improved SSDs such that conventional SSD elements, such as FTL tables, and processes, such as device-level garbage collection can be reduced or eliminated.

16. Memory devices are components used in computing systems to store and access data. A solid-state drive ("SSD") is a memory device that includes, among other components, a plurality of flash memory chips where the data is stored. In addition, an SSD contains a memory controller responsible for managing data access, read/write operations, and communications between the flash memory and a host. [1]

17. In the context of data storage, the host is typically a computing device, such as a computer or server, executing an operating system which accesses and uses the SSD for data storage and retrieval.[2]

18. SSDs are a worthwhile upgrade from their traditional mechanical counterparts.

19. An SSD will transform your computing experience with its revolutionary random access speeds, its multi-tasking prowess, and its outstanding durability and reliability.

20. A Hard Disk Drive ("HDD") works by way of a mechanical drive head that must physically move to access locations on a rapidly spinning magnetic disk.

21. An SSD has no moving parts and is therefore capable of accessing any location on the drive with equally fast speed and precision.

---

[1] *See generally* Ex. H, STORAGE NETWORKING INDUSTRY ASS'N, *What is an SSD?*, available at: https://www.snia.org/education/what-is-ssd..

[2] *Id.*; *see also* Ex. I, SAMSUNG, *Samsung Solid State Drive White Paper*, available at: https://download.semiconductor.samsung.com/resources/white-paper/Samsung_SSD_White_Paper.pdf, at 6, 7, 22, 24, 30 ("Samsung SSD White Paper").

22. SSDs are over 100 times faster than HDDs in accessing data.

23. SSDs deliver consistent performance thanks to their use of integrated circuits instead of physical spinning platters.

24. While rotating disks must wait for spindles, motors, heads, and arms to physically locate data locations, SSDs can access any location with lightning speed.

25. While rotating disks must wait for spindles, motors, heads, and arms to physically locate data locations, SSDs can access any location without having to wait for movement of physical parts.

26. An SSD benefits from massive latency advantages.

27. The ability to access any drive location without a performance penalty means you can have more applications open at the same time with less lag.

28. SSD users also experience dramatic improvements in boot time, shutdown, application loads, web browsing, application installations, and file copies.

29. Fast sequential speeds allow for quick file copies and smoother performance when working with large files, like videos.

30. An SSD can offer up to a 200x improvement in input/output operations per second ("IOPS") over a traditional HDD.

31. An SSD can withstand more than 40 times the vibration (20G) than an HDD (0.5G) and up to 1500G of shock (compared to less than 350Gs for a typical HDD).

32. The three most crucial components of any SSD are the NAND flash memory, the controller, and the firmware.

33. The NAND flash memory does the critical task of storing data.

34. The controller and the firmware work together to accomplish the complex and

essential tasks of managing data storage and maintaining the performance and lifespan of the SSD.

35. Using the control program contained in the drive firmware, the controller (an embedded microchip) executes automated signal processing, wear-leveling, error correcting code (ECC), bad block management, and garbage collection algorithms, communicates with the host device (e.g. a PC), and facilitates data encryption, among other tasks.

36. The controller also performs NAND maintenance such as scrubbing.

37. Redundantly stored in NAND flash, the firmware may be updated manually to improve or extend functionality when the SSD manufacturer releases an update.

38. The NAND flash itself is a kind of digital repository, while the controller functions as the command center for everything the SSD does – from actually reading and writing data to executing the garbage collection and wear-leveling algorithms that keep the drive clean and speedy.

39. NAND flash memory stores data in an array of memory cells made from floating-gate transistors.

40. Each cell has a finite lifetime, measured in terms of Program/Erase (P/E) cycles and affected by both process geometry (manufacturing technique) and the number of bits stored in each cell.

41. Data can be written into or read from NAND cells.

42. A write request or a write command is a request or command from the host OS to "write" or program data into the SSD.

43. A read request or a read command is a request or command from the host OS to "read" or access data from the SSD.

44. Physical NAND writes are not required to correspond directly to the space the host

system requests.

45. Utilizing a Flash Translation Layer ("FTL"), SSDs implement a logical to physical mapping system called Logical Block Addressing ("LBA").

46. The SSD assigns a physical address to each LBA written to the SSD by the host. The LBA can be used to locate the data regardless of the data's physical location. The FTL maintains a table containing the mapping of the LBAs to their physical locations in the underlying NAND.

47. FTLs require extra free space in an SSD.

48. FTLs often require significant DRAM to support their translation tables.

49. FTL tables require extra free space in an SSD.

50. The complexity of NAND storage necessitates some extra management processes, including bad block management, wear leveling, garbage collection (GC), and Error Correcting Code (ECC), all of which is managed by the device firmware through the SSD controller.

51. These management processes include NAND maintenance such as scrubbing.

52. The data stored in NAND flash is represented by electrical charges that are stored in each NAND cell.

53. Advanced wear-leveling code ensures that NAND cells wear out evenly (to prevent early drive failure and maintain consistent performance), while garbage collection algorithms preemptively prepare fresh storage space (to improve write performance).

54. Garbage collection algorithms preemptively prepare fresh storage space to consolidate valid data and make space for new writes.

55. Improved Error-Correcting Code (ECC) is able to detect and recover from errors at the bit level caused by the natural wear out of individual NAND cells.

56. Over-provisioning sets aside a certain portion of a conventional SSD's capacity exclusively to the controller.

57. Over-provisioning ensures that the SSD controller always has swap space available to accomplish its tasks.

58. As a precautionary measure, some vendors choose to pre-set a certain amount of mandatory over-provisioning at the factory.

59. There is the option to manually set aside additional space for even further improved performance (e.g. under demanding workloads).

60. The additional free space has the added benefit of providing the controller with unused capacity to use as a kind of non-official over-provisioning space.

61. Increased SSD capacity leads to increased performance and endurance.

62. Flash SSDs are capable of simultaneously accessing multiple NAND dies. This is referred to as parallelization.

63. NAND can be: single-level cell ("SLC"), which has 1 bit per cell; multi-level cell ("MLC"), which has 2 bits per cell; triple-level cell ("TLC"), which has 3 bits per cell; or quad-level cell ("QLC"), which has 4 bits per cell.

64. MLC offers higher data density than SLC. TLC offers higher data density than MLC. QLC offers higher data density than TLC.

65. Increasing the number of NAND chips in an SSD allows for more parallelization, helping to overcome the inherently slow program times that MLC NAND suffers compared to its SLC predecessor.

66. Samsung's fully integrated manufacturing approach means that it has intimate knowledge of every nuance of each component of an SSD.

67. Samsung's fully-integrated manufacturing approach means that it has intimate knowledge of every nuance of each component of its data center and enterprise SSDs.

68. In order to ensure that an SSD stays in prime working condition, the SSD controller must manage complex performance and lifetime optimization algorithms.

69. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty.

70. Each page belongs to a group of pages collectively known as a "block."

71. In order to write to a page that already contains data, it must first be erased.

72. While data is written in pages, it can only be erased in blocks.

73. An erase block can also be called an erase unit.

74. Hard disks are not subject to the same write/erase limitations that SSDs are – they can easily overwrite data in an existing location without erasing it first.

75. HDDs can overwrite data in an existing location without erasing it first.

76. Since SSD blocks cannot be overwritten, the SSD writes the new data to the next available page and simply marks the old data as "invalid."

77. Once a block is full of pages that all contain invalid data, that block is considered free and may be erased.

78. As an SSD is filled with more and more data, there will naturally be fewer free blocks readily available.

79. The SSD is then forced to actively consolidate valid data and prepare free blocks in order to write new data and perform maintenance.

80. This process of moving and consolidating data takes time, which is perceived as decreased performance, and requires free space.

81. This process of moving and consolidating data takes time, which results in decreased performance, and requires free space.

82. This is why over-provisioning, which guarantees a certain amount of free swap space to use for garbage collection and other maintenance activities, is so important for SSD performance – it allows the garbage collection algorithm to prepare free space in advance through data consolidation.

83. There are two methods of garbage collection, one that takes place in the background (when the SSD sits idle) and one that takes place in the foreground (when the SSD is working with data).

84. Samsung's SSDs actively manage data and create free blocks on demand.

85. TRIM is an optional command by which the operating system ("OS") can notify the SSD when data is no longer valid.

86. TRIM helps to make garbage collection more efficient by preparing invalid data for deletion.

87. In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level.

88. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit."

89. Error-correcting code ("ECC") uses this parity bit to compensate for other bits that may fail during normal operation of the drive.

90. When the controller detects a read failure, it will invoke ECC to try and recover from it.

91. If recovery is not possible, the firmware's bad block management feature will retire

the block and replace it with one of several free "reserved blocks."

92. "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

93. "Bad blocks" can be identified during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

94. "Bad blocks" can be marked during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

95. NAND flash memory suffers from the following limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called "P/E" cycles).

96. The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC).

97. In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells.

98. Using this technique, no single cell should be unduly stressed and prematurely fail.

99. If too many cells were to fail, the entire block would have to be retired.

100. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

101. Maintenance procedures like wear-leveling and garbage collection, which are created to deal with the unique properties of NAND flash memory, work together to help ensure that an SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan.

102. The key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability.

103. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component.

104. Over-provisioning, the practice of allocating a specific, permanent amount of free space on an SSD, is a widely used method for improving both SSD performance and endurance.

105. An SSD controller is responsible for the massive task of managing all data traffic and storage for the drive.

106. NAND technology's intrinsic complexities require a lot of extra work behind the scenes.

107. Each NAND cell has a limited lifespan.

108. Each NAND cell can only endure so many data reads/writes before it wears out or is no longer usable.

109. As NAND process technology advances, the chips themselves become increasingly smaller. As they shrink, the chips also become less reliable at holding data.

110. As NAND process technology advances, the topology of the chips shrinks, allowing for more space on a same-sized chip. As more space is fit into the same-sized chip, the chips also become less reliable at holding data.

111. Overwriting old data on NAND requires an erase of the entire NAND block (this same block may contain other data that is still valid).

112. The controller is constantly moving data around to ensure that the cells wear evenly and to preemptively prepare "free blocks" to use for future data writes.

113. Over-provisioning is a way to set aside a minimum amount of free space, inaccessible to the user or the OS, which the SSD controller can utilize as a kind of "work bench."

114. Over-provisioning can also be thought of as a pool, where only a percentage of

erase blocks can contain valid data. For instance, the pool may contain 120 erase blocks, but only a capacity of 100 erase blocks is advertised to the host, and most of the 120 erase blocks contain some valid data at any given moment. As erase blocks go bad, the pool shrinks to less than 120 erase blocks. Here, the 20 extra erase blocks in the pool over the 100 advertised to the host (20/100 or 20%) is the overprovisioning ratio.

115. A different over-provisioning ratio is recommended by usage applications and workload.

116. Activities that produce frequent read/write requests, especially random read/writes, put extra stress on the SSD, which in turn increases write amplification.

117. Write amplification refers to the phenomenon by which physical NAND writes outnumber logical write requests from the host.

118. Write amplification reflects the number of physical NAND writes for each logical write request or command sent by the host. Write amplification occurs because while a host might request that a piece of data be stored one time, the underlying SSD might copy (i.e., write) that same piece of data multiple times as it is managing the NAND flash, including through processes like garbage collection.

119. Heavy workloads wear out NAND cells faster, which increases the need for Bad Block management so that the controller can retire worn out cells.

120. Retiring a NAND cell requires the controller to copy all of the cell's valid data to a new block, taken from a limited number of "reserved" blocks.

121. Heavy workloads wear out NAND cells faster, which increases the need for Bad Block management so that the controller can retire worn out blocks.

122. Retiring a NAND block requires the controller to copy all of the block's valid data

to a new block, taken from a limited number of "reserved" blocks.

123. Working in tandem with bad block management, and reducing the need to use it altogether, is wear-leveling, which ensures that no single cell is written to more than others.

124. Working in tandem with bad block management is wear-leveling, which ensures that no single block is written to more than others.

125. To ensure free blocks are available, garbage collection algorithms consolidate good data and erase blocks of invalid data.

126. Over-provisioning helps with garbage collection, wear-leveling and bad block management by effectively increasing the size of the controller's "work bench," thus giving it extra free space to use while it consolidates and moves data or retires worn out cells.

127. Over-provisioning helps with garbage collection, wear-leveling and bad block management by effectively increasing the size of the controller's "work bench," thus giving it extra free space to use while it consolidates and moves data or retires worn out blocks.

128. Writing data to an SSD is not a simple one-to-one operation. There are complex processes taking place behind the scenes at all times.

129. Over-provisioning has a direct effect on SSD performance under sustained workloads and as the drive is filled with data.

130. Guaranteeing free space to accomplish the NAND management tasks like garbage collection, wear-leveling, and bad block management means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased, and recopied.

131. An added benefit is that over-provisioning makes all of the SSD maintenance procedures more efficient, reducing the write amplification factor ("WAF") by ensuring there's

room to work.

132. Samsung has complete control of both the NAND and the controller/firmware combination.

133. Samsung has complete control of both the NAND and the controller/firmware combination in its data center and enterprise products.

134. Having full control over the NAND, controller and firmware has major implications both for performance and reliability.

135. Samsung has intimate knowledge of every nuance of the NAND, controller, and firmware.

136. Samsung has intimate knowledge of every nuance of the NAND, controller, and firmware in its data center and enterprise products.

137. Samsung controls all of the most crucial design elements of an SSD: NAND, Controller, DRAM, and Firmware.

138. Samsung controls all of the most crucial design elements of its data center and enterprise SSDs: NAND, Controller, DRAM, and Firmware.

139. Samsung designs and manufactures every major SSD component.

140. Samsung designs and manufactures every major SSD component in its data center and enterprise SSDs.

C. Over-provisioning in Conventional SSDs

141. As with the previous section, this section provides background on over-provisioning in conventional SSDs. As a result of Radian's inventions, some of these features or elements can be reduced or eliminated such that they do not appear in the Accused Instrumentalities.

142. Over-provisioning is the technology that assigns a certain portion of a conventional SSD's capacity exclusively to the controller to allow the SSD to perform garbage collection more efficiently, helping to maintain steady state performance (i.e., sustained performance) and extend the SSD's endurance and lifetime.

143. The ratio of over-provisioning in conventional SSDs is calculated in percentages by dividing the over-provisioning capacity by the user capacity.

144. For example, in a total 512 GB SSD, if the over-provisioning is 112 GB and the user space is 400 GB, then its over-provisioning ratio is 28% (112 GB / 400 GB).

145. A data center operator or a hyperscaler may calculate the ratio of over-provisioning in conventional SSDs in percentages by dividing the sum of the over-provisioning capacity at the SSD level and the over-provisioning capacity at the data center or hyperscale system level by the user capacity.

146. For example, using that approach, in a total 512 GB SSD, if the over-provisioning is 112 GB at the SSD level, and the over-provisioning is an additional 88 GB at the data center or hyperscale system level, and the user space is 400 GB, then its over-provisioning ratio is 39% (200 GB / 512 GB).

147. Under a heavy workload (for example, a server, data center or heavy workload client PC applications), a minimum of 6.7% over-provisioning is recommended and over 20% and even 50% has been used.

148. Activities that produce frequent read/write requests, especially random read/writes, put extra stress on the SSD, which in turn increases write amplification (a phenomenon by which physical NAND writes outnumber logical write requests from the host).

149. Over-provisioning has a direct effect on the SSD's random performance as the drive

is filled with data.

150. Guaranteeing free space to accomplish the NAND management tasks (garbage collection, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied.

151. An added benefit is that over-provisioning makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor by ensuring there is room to work, which improves the SSD's lifetime.

152. An SSD with 20% OP would have a much longer lifetime, and have better performance, than an SSD with 10% OP.

153. Some SSDs are already over-provisioned by the manufacturer and users cannot access or control the over-provisioning ratio.

154. Samsung has made SSDs which are already over-provisioned.

155. Samsung has made SSDs for which users cannot access or control the pre-set over-provisioning ratio.

156. Samsung has made SSDs for data centers with an over-provisioning ratio of 6.7%.

157. Samsung has made SSDs for data centers with an over-provisioning ratio of 15%.

158. Samsung has made SSDs for data centers with an over-provisioning ratio of 20%.

159. Samsung has made SSDs for data centers with an over-provisioning ratio of 28%.

160. Samsung has made SSDs for data centers with an over-provisioning ratio of 33%.

161. Samsung has made SSDs for data centers with an over-provisioning ratio of 35%.

162. Samsung has sold SSDs for data centers with an over-provisioning ratio of 6.7%.

163. Samsung has sold SSDs for data centers with an over-provisioning ratio of 15%.

164. Samsung has sold SSDs for data centers with an over-provisioning ratio of 20%.

165. Samsung has sold SSDs for data centers with an over-provisioning ratio of 28%.

166. Samsung has sold SSDs for data centers with an over-provisioning ratio of 33%.

167. Samsung has sold SSDs for data centers with an over-provisioning ratio of 35%.

168. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 6.7%.

169. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 15%.

170. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 20%.

171. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 28%.

172. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 33%.

173. Samsung has marketed SSDs for data centers with an over-provisioning ratio of 35%.

D. Performance Measurements of SSDs

174. SSD capacities are marketed in megabytes ("MB"), gigabytes ("GB"), terabytes ("TB"), petabytes ("PB"), zetabytes ("ZB"), and so forth. These annotations suggest a capacity that is in decimal form (i.e., 10^x), such that, for example, 1 MB indicates 1,000,000 bytes (or 10^6 bytes). This convention can be used when generally referring to data storage capacity, such as in marketing or when describing or discussing SSD capacities in general. When specifying the capacities precisely in a technical sense, SSD capacities are denoted as mebibytes ("MiB"),

gibibytes ("GiB"), tebibytes ("TiB"), pebibytes ("PiB"), zebibytes ("ZiB"), and so forth . These annotations specify a capacity that is in binary form (i.e., 2^x), such that, for example, 1 MiB means 1,048,576 bytes (or 2^{20} bytes). This convention can be used when precise calculations are important.

175. Sequential read speeds are reported in megabytes per second (MB/s) and indicate how fast the SSD will be at completing tasks like accessing applications using large sequential files.

176. Sequential read speeds are reported in megabytes per second (MB/s) and indicate how fast the SSD will be at completing tasks like accessing large multimedia files, transcoding, game level loading, some types of game play, watching and editing video.

177. Sequential read is the speed at which the drive can read data from contiguous memory spaces.

178. Sequential write speeds, also reported in megabytes per second (MB/s), indicate how fast the SSD will be at tasks like writing or storing large sequential files.

179. Sequential write speeds, also reported in megabytes per second (MB/s), indicate how fast the SSD will be at tasks like application installation and document backup.

180. Sequential write is the speed at which the drive can write data to contiguous memory spaces.

181. Random read speeds, reported in input/output operations per second ("IOPS"), indicate how fast the SSD will be at completing tasks that read small data structures from random locations.

182. Random read speeds, reported in input/output operations per second ("IOPS"), indicate how fast the SSD will be at completing tasks like antivirus scans, searching for email in

Outlook, web browsing, application loading, PC booting, or working in Microsoft Word.

183. Random read is the speed at which the drive can read data from non-contiguous memory spaces.

184. Random write speeds, also reported in IOPS, indicate how fast the drive will be able to complete tasks that write small data structures to random locations.

185. Random write speeds, also reported in IOPS, indicate how fast the drive will be able to complete tasks like downloading email, compressing files, or browsing the web.

186. Random write is the speed at which the drive can write data to non-contiguous memory spaces.

187. Idle time is the period of time when an SSD is not handling requests to read or write data.

188. Latency is the response time of the SSD, or the time between when you initiate a request to read or write data and when the SSD completes the request.

189. Latency is the response time of the SSD, or the time between when you initiate a request to read or write data and when the SSD starts to complete the request.

190. Latency is reported in microseconds (µs) or milliseconds (ms).

191. Queue depth/concurrent IOs is the outstanding (pending) number of IO requests that must be handled by the drive.

192. Queue depth/concurrent IOs is the outstanding (pending) number of IO requests that are issued to the drive.

193. Bandwidth is the amount of data that can be transferred in a given time.

194. Bandwidth is reported in megabytes per second (MB/s) or gigabytes per second (GB/s).

E. Zoned Namespaces ("ZNS")

195. In March 2018, at the Open Compute Project ("OCP") Summit, Microsoft gave a presentation titled, "Denali: The Next-Generation High-Density Storage Interface." This presentation discussed the need for improving the efficiency of flash in enterprise and cloud-based applications and outlined concepts for a new type of storage interface optimized to do so. With this presentation, Microsoft announced Project Denali, which was expected to be a new standard for SSDs designed with cloud-based applications in mind, and which would strive to maximize flexibility. Led by Microsoft, this OCP initiative led to the creation of a working group within the NVM Express ("NVMe") organization that later developed the ZNS specification.

196. In 2020, NVMe introduced the NVMe Zoned Namespace ("ZNS") Command Set Specification.[3] The ZNS Specification describes an architecture for SSDs that divides the memory space into distinct zones that are accessible to the host and are to be erased as a whole.[4] It was promoted as being able to greatly increase storage efficiency, reduce write amplification, optimize read/write performance, and extend the lifespan of an SSD.[5]

197. The NVM Express Base Specification defines an interface for host software to communicate with a non-volatile memory subsystem (NVM subsystem).

---

[3] *See* Ex. J, Matias Bjørling, *New NVMe™ Specification Defines Zoned Namespaces (ZNS) as Go-To Industry Technology*, NVM EXPRESS, available at: https://nvmexpress.org/new-nvmetm-specification-defines-zoned-namespaces-zns-as-go-to-industry-technology/; NVM EXPRESS, *NVMe Zoned Namespaces (ZNS) Command Set Specification*, available at https://nvmexpress.org/specification/nvme-zoned-namespaces-zns-command-set-specification/.

[4] *See* Ex. K, Billy Tallis, *The Next Step in SSD Evolution: NVMe Zoned Namespaces Explained*, available at: https://www.anandtech.com/print/15959/nvme-zoned-namespaces-explained; NVM EXPRESS, *NVMe Zoned Namespaces (ZNS) Command Set Specification*, available at https://nvmexpress.org/specification/nvme-zoned-namespaces-zns-command-set-specification/.

[5] *Id*.

198. The Zoned Namespace Command Set Specification defines additional functionality for the Zoned Namespace Command Set.

199. The Zoned Namespace Command Set is based on the NVM Command Set (which is defined in the NVM Command Set Specification).

200. The Command Set Identifier (CSI) value for the Zoned Namespace Command Set is 02h.

201. A namespace is a set of resources that may be accessed by a host and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by the Zoned Namespace Command Set Specification.

202. A namespace has an associated namespace identifier that a host uses to access that namespace.

203. A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set.

204. A zoned namespace is divided into a set of equally sized zones, which are contiguous non-overlapping ranges of logical block addresses.

205. Each zone has an associated Zone Descriptor that contains a set of attributes.

206. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

207. The host may use the Zone Management Receive command to determine the current write pointer for a zone.

208. The Zone Management Receive command returns a data buffer that contains information about zones. That information includes characteristics of the zone, the state of the zone, the capacity of the zone, and other information defined in the Zoned Namespace Command

Set Specification. The host uses this command to determine the current settings for this information.

209. The Data Pointer field specifies the location of a data buffer where data is transferred from.

210. The Starting LBA field specifies an LBA in the lowest numbered zone that the Zone Receive Action operates on.

211. The Zone Management Receive command Zone Receive Action field specifies what action to perform.

212. The Report Zones action returns the Report Zones data structure.

213. The Zone Descriptors of the Report Zones data structure shall: a) report only Zone Descriptors of zones for which the ZSLBA value is greater than or equal to the ZSLBA value of the zone specified by the SLBA value in the command; b) match the criteria in the Zone Receive Action Specific field; and c) be sorted in ascending order by the ZSLBA value of each zone.

214. The Extended Report Zones action returns the Extended Report Zones data structure.

215. The Zone Descriptors and Zone Descriptor Extensions of the Extended Report Zones data structure shall: a) report only Zone Descriptors and Zone Descriptor Extensions of zones for which the ZSLBA value of is greater than or equal to the ZSLBA value of the zone specified by the SLBA value in the command; b) match the criteria in the Zone Receive Action Specific field; and c) be sorted in ascending order by the ZSLBA value of each zone.

216. Reset Zone Recommended is a Zone Attribute.

217. If the Reset Zone Recommended bit is set to “1”, then the controller recommends that this zone be reset.

218. Reset Zone Recommended Time Limit is Zone Attribute Information that indicates additional information associated with the Reset Zone Recommended attribute for the zone.

219. If the Reset Zone Recommended bit is set to "1", then the value in the Reset Zone Recommended Time Limit field selects a field in the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended bit is set to "1" in the Zone Attributes field for that zone.

220. The Zone Starting Logical Block Address ("ZSLBA") field contains the 64-bit address of the lowest logical block for the zone.

221. The Write Pointer field is the logical block address where the next write operation for this zone should be issued.

222. The Zone Management Send command requests an action on one or more zones.

223. The Zone Management Send command with a Zone Send Action of Reset Zone sets the write pointer to the ZSLBA for that zone.

F. Radian's Innovations

224. Radian's founders began working in 2010 on solutions and technology to solve the inefficiencies of flash management in data center storage applications, specifically by reducing overhead and write amplification.

225. Radian is responsible for the creation and development of Cooperative Flash Management ("CFM"), a novel system-driven approach to flash management.

226. Previously, flash memory management could be performed via the use of a flash translation layer ("FTL"), a system that, while reliable, suffered from general inefficiencies. Later developments gave rise to software-defined flash ("SDF"). Basic SDF devices, while configurable

for more efficient operations, require massive development efforts and are significantly more complex to integrate and maintain than a FTL SSD implementation.

227. CFM, the flash management technology developed by Radian, redistributes flash management responsibilities between the host and the device. Based on a revolutionary new SSD firmware architecture, host-side capabilities can be leveraged to eliminate redundant functionality, such as different levels of space management, and redundant processes such as device-level garbage collection, to offer dramatically higher performance and more efficient utilization of flash storage.

228. Radian frequently presented its innovative approach to the rest of the industry.

229. For example, Radian gave a presentation at the Flash Memory Summit on replacing the flash translation layer ("FTL"), the conventional approach to managing flash memory, with CFM. Among other virtues of CFM, this presentation stressed how CFM provided a "Goldilocks" type alternative to the extremes of FTL and "true" SDF, offering the advantages of both without their respective downsides.

230. Later that year, at the 2015 SNIA Storage Developer Conference (SDC), Radian gave a presentation on the integration of CFM in data center or hyperscale systems.

231. On August 11, 2015, Radian presented its Symphonic CFM at the Flash Memory Summit, where it won Most Innovative Software Technology. At this presentation Radian described the advantages and relative trade-offs between different flash management solutions.

232. Additionally, on September 21, 2016, Radian gave a presentation on software-defined flash, and the utility of CFM, at the SNIA Storage Developer Conference.

233. Radian would go on to exhibit at FMS every year from 2015 through 2019.

234. On August 14, 2019, Radian submitted an entry again to the Flash Memory Summit

for its Zoned Flash technology. At that conference, Radian gave a presentation on its implementation of Zoned Flash, building on the prior success of its Symphonic CFM, and publicly demonstrating Zoned Flash operating in a system—an industry first—and under industry standard performance benchmark measurements. Radian's implementation of Zoned Namespaces firmware won Most Innovative Technology in the SSD Firmware category at this FMS 2019.

235. Soon after, on September 18, 2019, at the SNIA Storage Developer Conference, Radian presented its Zoned Flash technology and demonstrated Zoned Flash SSDs operating in Dual Port mode and Multi-Drive Volumes—another industry first.

236. Radian did not just give presentations, either; it embraced head-to-head comparisons between its products and other solutions in the field. Radian also supplied functional evaluation units under confidentiality agreements for review and testing by potential customers or partners.

237. On August 7, 2019, at the Flash Memory Summit, Radian demonstrated a benchmark comparison between its RMS-350 NVMe SSD with its Cooperative Zoned Flash implementation and a production-released FTL-based NVMe SSD.

238. Additionally, in February 2020, Radian provided a live demonstration of IBM's software-defined log-structured array storage management stack using Radian's Zoned Flash SSD for the first time.

239. This demonstration proved the benefits of Radian's technology over existing state-of-the-art SSDs. Compared to a near-identical FTL-based SSD with the same configurations, IBM Research observed block level improvements of 3x in throughput, 50x in tail latencies, and an estimated 3x improvement in wear out, as well as system-level improvements of 65% in transactions per second ("tps"), over 22x in tail latencies, and an estimated 3x improvement in

wear out.

240. However, Radian's advances in the field were also not confined to one-off demonstrations.

241. On July 21, 2020, Radian announced that, in an industry first, it had integrated its RMS-350 U.2 Zoned Flash SSD with the Storage Performance Development Kit (SPDK), a set of tools and libraries for writing high-performance storage applications. At the same time, Radian announced that it had successfully qualified the RMS-350 under the JESD218 testing requirements.

242. On September 2, 2020, Radian announced that its RMS-350 ZNS SSD passed compliance testing with a third-party test suite for the NVMe ZNS Command Set, the first in the industry to do so, and merely months after the first working draft of the ZNS specification was released by NVMe.

243. This was despite the fact that Radian was never provided with a working draft of the ZNS specification prior to the publication of the final specification.

G. Radian and Samsung

244. Radian first reached out to Samsung in 2015 to discuss its patented-protected technology.

245. Since 2015, Radian has discussed its technology with a number of individuals at Samsung.

246. At the August 2015 Flash Memory Summit, Radian gave a presentation titled, "Replacing the FTL with Cooperative Flash Management."

247. Conference participants attended Radian's presentation at the August 2015 Flash Memory Summit.

248. Conference participants visited Radian's exhibit booth at the August 2015 Flash Memory Summit.

249. Conference participants spoke and interacted with Radian representatives at the August 2015 Flash Memory Summit.

250. Samsung employees attended the August 2015 Flash Memory Summit.

251. Samsung employees who attended the August 2015 Flash Memory Summit did so within the scope of their employment at Samsung.

252. At the September 2015 SNIA Storage Developer Conference, Radian gave a presentation titled, "Integrating Cooperative Flash Management with SMR Technology for Optimized Tiering in Hybrid Systems."

253. Conference participants attended Radian's presentation at the September 2015 SNIA Storage Developer Conference.

254. Conference participants visited Radian's exhibit booth at the September 2015 SNIA Storage Developer Conference.

255. Conference participants spoke and interacted with Radian representatives at the September 2015 SNIA Storage Developer Conference.

256. Samsung employees attended the September 2015 SNIA Storage Developer Conference.

257. Samsung employees who attended the September 2015 SNIA Storage Developer Conference did so within the scope of their employment at Samsung.

258. At the September 2016 SNIA Storage Developer Conference, Radian gave a presentation titled, "Software-Defined Flash: Trade-Offs of FTL, Open-Channel, and Cooperative Flash Management."

259. Conference participants attended Radian's presentation at the September 2016 SNIA Storage Developer Conference.

260. Conference participants visited Radian's exhibit booth at the September 2016 SNIA Storage Developer Conference.

261. Conference participants spoke and interacted with Radian representatives at the September 2016 SNIA Storage Developer Conference.

262. Samsung employees attended the September 2016 SNIA Storage Developer Conference.

263. Samsung employees who attended the September 2016 SNIA Storage Developer Conference did so within the scope of their employment at Samsung.

264. At the August 2019 Flash Memory Summit, Radian gave a presentation titled, "Fast Integration and Furious Performance with Zoned Flash Drives."

265. Conference participants attended Radian's presentation at the August 2019 Flash Memory Summit.

266. Conference participants visited Radian's exhibit booth at the August 2019 Flash Memory Summit.

267. Conference participants spoke and interacted with Radian representatives at the August 2019 Flash Memory Summit.

268. Samsung employees attended the August 2019 Flash Memory Summit.

269. Samsung employees who attended the August 2019 Flash Memory Summit did so within the scope of their employment at Samsung.

270. At the September 2019 SNIA Storage Developer Conference, Radian gave a presentation titled, "Zoned Flash SSDs in Advanced Storage Systems."

271. Conference participants attended Radian's presentation at the September 2019 SNIA Storage Developer Conference.

272. Conference participants visited Radian's exhibit booth at the September 2019 SNIA Storage Developer Conference.

273. Conference participants spoke and interacted with Radian representatives at the September 2019 SNIA Storage Developer Conference.

274. Samsung employees attended the September 2019 SNIA Storage Developer Conference.

275. Samsung employees who attended the September 2019 SNIA Storage Developer Conference did so within the scope of their employment at Samsung.

276. Samsung is a member of the NVM Express ("NVMe") organization.

277. Samsung is a Promoter-level member of NVMe.

278. In NVMe, the Promoter level of membership is the top tier of membership.

279. Samsung employees represent Samsung at NVMe meetings.

280. When Samsung employees attend NVMe meetings, they do so within the scope of their employment at Samsung.

281. Samsung employees collaborate with employees of other companies that are members of NVMe.

282. When Samsung employees collaborate with employees of other companies that are members of NVMe, they do so within the scope of their employment at Samsung.

283. Samsung first learned of Radian in 2010.

284. Samsung was a member of NVMe at the time it learned of Radian.

285. Samsung became a member of NVMe after it learned of Radian.

286. Samsung was a member of NVMe when ZNS specification was first published.

287. Samsung employees were authors on the ZNS specification.

288. When Samsung employees authored portions of the ZNS specification, they did so within the scope of their employment at Samsung.

289. Samsung employees contributed material to the ZNS specification.

290. When Samsung employees contributed material to the ZNS specification, they did so within the scope of their employment at Samsung.

291. Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian.

292. When Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian, they did so within the scope of their employment.

293. Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian's technology.

294. When Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian's technology, they did so within the scope of their employment.

295. Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian's patented technology.

296. When Samsung employees contributed materials to the ZNS specification after Samsung learned about Radian's patented technology, they did so within the scope of their employment.

297. Employees of companies that were NVMe members tried to pressure Radian to become a member of NVMe.

298. Executives of companies that were NVMe members tried to pressure Radian to

become a member of NVMe.

299. NVMe requires that, effective upon NVMe's adoption of a final specification, its members grant to each other member a non-exclusive, non-transeferable, non-sublicensable, world-wide license in, to, and under any of that member's patent claims that are necessary to implement a compliant portion of that specification. Members are also required to grant this license on a royalty free basis.

300. NVMe's policies state that, effective upon NVMe's adoption of a final specification, its members must grant to each other member a non-exclusive, non-transeferable, non-sublicensable, world-wide license in, to, and under any of that member's patent claims that are necessary to implement a compliant portion of that specification. NVMe's policies also state that members must grant this license on a royalty free basis.

301. NVMe did not provide Radian with any working drafts of the ZNS specification before it was publicly released.

302. Radian developed a ZNS SSD before any NVMe member company did.

303. Samsung knew that Radian developed a ZNS SSD before any NVMe Member company did.

304. Radian developed an SSD that complied with the ZNS specification before any NVMe Member company did.

305. Samsung knew that Radian developed an SSD that complied with the ZNS specification before any NVMe Member company did.

306. Radian developed an SSD that was certified by an independent lab as complying with the ZNS specification before any NVMe Member company did.

307. Samsung knew that Radian developed an SSD that was certified by an independent

lab as complying with the ZNS specification before any NVMe Member company did.

308. Radian developed a ZNS SSD before Samsung did.

309. Samsung knew that Radian developed a ZNS SSD before Samsung did.

310. Radian developed an SSD that complied with the ZNS specification before Samsung did.

311. Samsung knew that Radian developed an SSD that complied with the ZNS specification before Samsung did.

312. Radian developed an SSD that was certified by an independent lab as complying with the ZNS specification before Samsung did.

313. Samsung knew that Radian developed an SSD that was certified by an independent lab as complying with the ZNS specification before Samsung did.

314. Radian is not a member of the NVMe.

315. Radian has never been a member of NVMe.

316. Samsung has not taken a license to practice Radian's patents.

317. Samsung has no implied license to practice Radian's patents.

318. Samsung has no right to practice Radian's patents.

319. Samsung has no authority to practice Radian's patents.

H. Samsung and ZNS

320. On June 2, 2021, Samsung issued a press release announcing the debut of the

PM1731a, its first ZNS SSD. [6]

321. In the press release, Samsung included quotes from Dr. Sangyeun Cho.

322. Dr. Cho was the Senior Vice President of the Memory Software Development Team at Samsung Electronics.

323. According to Samsung, this new SSD leveraging ZNS technology "will maximize available user capacity and offer an extended lifespan in storage server, data center and cloud environments."[7]

324. In its press release, Samsung also discussed the main benefits of ZNS, which include achieving a write amplification factor close to 1, enabling the device to last up to four times longer than conventional NVMe SSDs, and eliminating the need for overprovisioning to free up more memory space for data storage.[8]

325. Samsung engaged in a joint effort with Western Digital to establish the Zoned Storage Technical Work Group in December 2021, and signed a memorandum with Western Digital in March 2022 for a collaboration to standardize and drive the adoption of Zoned Storage solutions, which include ZNS SSD solutions.[9]

---

[6] *See* Ex. L, SAMSUNG, *Samsung Introduces Its First ZNS SSD With Maximized User Capacity and Enhanced Lifespan*, available at: https://semiconductor.samsung.com/news-events/news/samsung-introduces-its-first-zns-ssd-with-maximized-user-capacity-and-enhanced-lifespan/.

[7] *Id*.

[8] *Id*.

[9] *See* Ex. M, SAMSUNG, *Samsung and Western Digital Begin Far-Reaching Collaboration To Drive Standardization of Next-generation Storage Technologies for Broader Ecosystem Support and Customer Adoption*, available at: https://news.samsung.com/global/samsung-and-western-digital-begin-far-reaching-collaboration-to-drive-standardization-of-next-generation-storage-technologies-for-broader-ecosystem-support-and-customer-adoption

326. Samsung has also been promoting the ZNS technology in industry meetings. For example, during the September 2020 SNIA Storage Developer Conference, Samsung's principal software engineer Javier González and staff engineer Kanchan Joshi gave a presentation on behalf of Samsung explaining ZNS and its use cases.[10] In the same year, in another industry conference called Open Compute Project, Javier González gave another presentation that introduced the benefits of ZNS on behalf of Samsung.[11] During the 2022 SDC, Samsung gave a presentation titled, "A New Adapter for Zoned Namespace SSDs" through its engineer Hui Qi.[12] In that meeting, Samsung explained the benefits of ZNS by comparing a ZNS SSD with a conventional SSD, and called for efforts to "build[] ZNS ecosystem together."[13]

327. Samsung has incorporated ZNS technology into a significant portion of its SSD products (the "Accused Instrumentalities"). Samsung's Accused Instrumentalities include the

---

[10] Ex. N, Samsung, *ZNS: Enabling In-place Updates and Transparent High Queue-depths*, STORAGE NETWORKING INDUSTRY ASS'N, available at: https://snia.org/sites/default/files/SDC/2020/050-Gonz%C3%A1lez-Joshi-ZNS-Enabling-In-place-Updates.pdf

[11] Ex. O, Sched, *Zoned Namespaces: From Archival to I/O Predictability - presented by Samsung*, available at: https://2020ocpvirtualsummit.sched.com/event/bYb0/on-demand-zoned-namespaces-from-archival-to-io-predictability-presented-by-samsung.

[12] Ex. P, Hui Qi, *A New Adapter for Zoned Namespace SSDs*, available at: https://www.sniadeveloper.org/sites/default/files/SDC/2022/pdfs/SNIA-SDC22-Hui-A-New-Adapter-for-ZNS-SSD.pdf

[13] *Id*.

PM1731a, PM1733,[14] PM1735,[15] PM9731a,[16] and PB SSD[17] products, as well as other Samsung SSD products that implement ZNS, including any products rebranded for other companies or made directly for partners or other companies under a white label or private label.[18] Additionally, Accused Instrumentalities include other Samsung SSD products that implement functionalities substantially similar to those in ZNS, even if not marketed as a ZNS SSD or an SSD that implements ZNS or complies with the ZNS specification.

328. The Accused Instrumentalities make use of commands defined in a NVMe specification.

329. The Accused Instrumentalities receive commands defined in a NVMe specification.

330. The Accused Instrumentalities respond to commands defined in a NVMe specification.

331. The Accused Instrumentalities are capable of using commands defined in a NVMe

---

[14] *See* Ex. Q, SAMSUNG, *Samsung PM1733 NVMe SSD Product Brief*, available at: https://download.semiconductor.samsung.com/resources/brochure/PM1733%20NVMe%20SSD.pdf (the "Samsung PM1733 Product Brief").

[15] Samsung markets the PM1735 with the PM1733. *See* Ex. R, SAMSUNG, *PM1733/PM1735 Enterprise SSD*, available at: https://semiconductor.samsung.com/us/ssd/enterprise-ssd/pm1733-pm1735/.

[16] *See* Ex. S, Jonghyeok Park, *VLDB Lab ZNS Guide Samsung ZNS SSD*, GITHUB, available at: https://github.com/JonghyeokPark/ZNS.

[17] *See* Ex. T, Patrick Kennedy, *Samsung PB SSD at 128TB for 60PB All-flash Racks*, available at: https://www.servethehome.com/samsung-pb-ssd-at-128tb-for-60pb-all-flash-racks-intel-inspur/; https://semiconductor.samsung.com/emea/news-events/tech-blog/samsung-electronics-presents-next-generation-memory-solutions-to-settle-challenges-in-the-big-data-era-at-fms-2022/ (describing PB SSD and other products being presented at the 2022 Flash Memory Summit).

[18] *See, e.g.*, Ex. U, SUPPLY SERVER, *HPE PM1733 P16456-002 SSD 3.84 TB PCIe 4.0 x4 (NVMe) Refurbished*, available at: https://www.serversupply.com/SSD%20W-TRAY/NVMe/3.84TB/HPE/P16456-002_348219.htm; *see also* https://smrc.biz.samsung.com/about-us (describing partners such as NetApp in relation to ZNS SSDs).

specification.

332. The Accused Instrumentalities are capable of receiving commands defined in a NVMe specification.

333. The Accused Instrumentalities are capable of responding to commands defined in a NVMe specification.

334. The Accused Instrumentalities make use of Identify commands.

335. The Accused Instrumentalities receive Identify commands.

336. The Accused Instrumentalities respond to Identify commands.

337. The Accused Instrumentalities are capable of using Identify commands.

338. The Accused Instrumentalities are capable of receiving Identify commands.

339. The Accused Instrumentalities are capable of responding to Identify commands.

340. The Accused Instrumentalities make use of Write commands.

341. The Accused Instrumentalities receive Write commands.

342. The Accused Instrumentalities respond to Write commands.

343. The Accused Instrumentalities are capable of using Write commands.

344. The Accused Instrumentalities are capable of receiving Write commands.

345. The Accused Instrumentalities are capable of responding to Write commands.

346. The Accused Instrumentalities make use of Read commands.

347. The Accused Instrumentalities receive Read commands.

348. The Accused Instrumentalities respond to Read commands.

349. The Accused Instrumentalities are capable of using Read commands.

350. The Accused Instrumentalities are capable of receiving Read commands.

351. The Accused Instrumentalities are capable of responding to Read commands.

352. The Accused Instrumentalities make use of Zone Management Receive commands.

353. The Accused Instrumentalities receive Zone Management Receive commands.

354. The Accused Instrumentalities respond to Zone Management Receive commands.

355. The Accused Instrumentalities are capable of using Zone Management Receive commands.

356. The Accused Instrumentalities are capable of receiving Zone Management Receive commands.

357. The Accused Instrumentalities are capable of responding to Zone Management Receive commands.

358. The Accused Instrumentalities make use of Zone Management Send commands.

359. The Accused Instrumentalities receive Zone Management Send commands.

360. The Accused Instrumentalities respond to Zone Management Send commands.

361. The Accused Instrumentalities are capable of using Zone Management Send commands.

362. The Accused Instrumentalities are capable of receiving Zone Management Send commands.

363. The Accused Instrumentalities make use of Zone Append commands.

364. The Accused Instrumentalities receive Zone Append commands.

365. The Accused Instrumentalities respond to Zone Append commands.

366. The Accused Instrumentalities are capable of using Zone Append commands.

367. The Accused Instrumentalities are capable of receiving Zone Append commands.

368. The Accused Instrumentalities are capable of responding to Zone Append commands.

369. The Accused Instrumentalities make use of Report Zones data structures.

370. The Accused Instrumentalities send Report Zones data structures.

371. The Accused Instrumentalities are capable of using Report Zones data structures.

372. The Accused Instrumentalities are capable of sending Report Zones data structures.

373. The Accused Instrumentalities perform Reset Zone.

374. The Accused Instrumentalities are capable of performing Reset Zone.

375. The Accused Instrumentalities erase the erase units (or erase blocks) in a zone.

376. The Accused Instrumentalities are capable of erasing the erase units (or erase blocks) in a zone.

377. The Accused Instrumentalities update the write pointer in a zone.

378. The Accused Instrumentalities are capable of updating the write pointer in a zone.

379. The Accused Instrumentalities transition the state of a zone.

380. The Accused Instrumentalities are capable of transitioning the state of a zone.

I. Radian's Prevalence in Samsung's Patent Prosecution

381. Radian patent applications, patent application publications, and issued patents have been cited in the prosecution of numerous Samsung patent applications, including both citations by examiners against Samsung as well as being identified and cited by Samsung itself.

382. USPTO examiners have cited a Radian patent application, patent application publication, or issued patent against multiple Samsung patent applications.

383. Samsung has cited a Radian patent application, patent application publication, or issued patent in the prosecution of multiple Samsung patent applications.

384. On February 29, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/055,689 ("Samsung's '689 Application"), entitled "Storage device and operating method of

storage device," which issued on April 16, 2019, as U.S. Pat. No. 10,261,697 ("Samsung's '697 Patent").

385. During the prosecution of Samsung's '697 Patent, the examiner cited Radian's U.S. Pat. No. 9,652,376 ("Radian's '376 Patent") against Samsung's '689 Application.

386. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '689 Application.

387. On April 5, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/090,799 ("Samsung's '799 Application"), entitled "Heuristic interface for enabling a computer device to utilize data property-based data placement inside a nonvolatile memory device," which issued on December 17, 2019, as U.S. Pat. No. 10,509,770 ("Samsung's '770 Patent").

388. During the prosecution of Samsung's '770 Patent, Samsung Electronics Co., Ltd. cited Radian's U.S. Pat. Pub. No. 2014/0215129 ("the '129 Publication"), which was the published application of Radian's U.S. Pat. App. No. 13/767,723 ("the '723 Application"), which later issued as Radian's '376 Patent.

389. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

390. On May 2, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/144,588 ("Samsung's '588 Application"), entitled "Data property-based data placement in a nonvolatile memory device," which issued on October 4, 2022, as U.S. Pat. No. 11,461,010 ("Samsung's '010 Patent").

391. During the prosecution of Samsung's '010 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

392. Radian's '129 Publication was the published application of Radian's '723

Application, which later issued as Radian's '376 Patent.

393. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

394. On November 4, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/344,422 ("Samsung's '422 Application"), entitled "Smart I/O stream detection based on multiple attributes," which issued on May 7, 2019, as U.S. Pat. No. 10,282,324 ("Samsung's '324 Patent").

395. During the prosecution of Samsung's '324 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

396. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

397. On April 4, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 16/844,995 ("Samsung's '995 Application"), entitled "Self-configuring SSD multi-protocol support in host-less environment," which issued on May 14, 2024, as U.S. Pat. No. 11,983,138 ("Samsung's '138 Patent").

398. During the prosecution of Samsung's '138 Patent, Samsung Electronics Co., Ltd. cited Radian's U.S. Pat. No. 9,400,749 ("the '749 Patent").

399. Radian's '749 Patent, which was cited by Samsung Electronics Co., Ltd. is a continuation of Radian's '376 Patent, from which the '183 Patent asserted by Radian in this action is descended.

400. On October 27, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/336,772 ("Samsung's '772 Application"), entitled "IO redirection methods with cost estimation." The '772 Application was published on May 4, 2017, as U.S. Pat. Pub. No.

2017/0123700, and was ultimately abandoned.

401. During the prosecution of Samsung's '772 Application, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

402. During the prosecution of Samsung's '772 Application, Samsung Electronics Co., Ltd. also cited Radian's U.S. Pat. No. 9,229,854 ("the '854 Patent").

403. Radian's '854 Patent, which was cited by Samsung Electronics Co., Ltd. is a descendant of Radian's '376 Patent.

404. The '183 Patent asserted by Radian in this action is also a descendant of Radian's '376 Patent, which was also cited by Samsung Electronics Co., Ltd.

405. On February 17, 2017, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/046,435 ("Samsung's '435 Application"), entitled "Coordinated garbage collection of flash devices in a distributed storage system," which issued on April 9, 2019, as U.S. Pat. No. 10,254,998 ("Samsung's '998 Patent").

406. During the prosecution of Samsung's '998 Patent, the examiner cited Radian's '129 Publication and Radian's '854 Patent against Samsung's '435 Application.

407. Radian's '129 Publication was the published application of Radian's '376 Patent, and Radian's '854 Patent, which was cited against Samsung Electronics Co., Ltd., is a continuation of Radian's '376 Patent.

408. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was also cited against Samsung's '435 Application.

409. On January 19, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/001,217 ("Samsung's '217 Application"), entitled "Mitigating GC effect in a raid configuration," which issued on October 31, 2017, as U.S. Pat. No. 9,804,787 ("Samsung's '787

Patent").

410. During the prosecution of Samsung's '787 Patent, the examiner cited Radian's '129 Publication and Radian's '854 Patent.

411. Radian's '129 Publication, which was cited against Samsung's '435 Application, was the published application of Radian's '376 Patent.

412. Radian's '854 Patent, which was cited against Samsung's '435 Application, is a continuation of Radian's '376 Patent.

413. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited against Samsung's '435 Application.

414. On November 20, 2015, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 14/947,931 ("Samsung's '931 Application"), entitled "Distributed multimode storage management," which issued on April 17, 2018, as U.S. Pat. No. 9,946,642 ("Samsung's '642 Patent").

415. During the prosecution of Samsung's '642 Patent, Samsung Electronics Co., Ltd. cited Radian's U.S. Pat. Pub. No. 2014/0365719 ("the '719 Publication").

416. Radian's '719 Publication, which was cited by Samsung Electronics Co., Ltd., was the published application of Radian's U.S. Pat. App. No. 14/466,167, which issued as Radian's '454 Patent.

417. The '183 Patent asserted by Radian in this action is a descendant of Radian's '454 Patent.

418. On November 13, 2015, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 14/941,512 ("Samsung's '512 Application"), entitled "Multimode storage management system," which issued on June 4, 2018, as U.S. Pat. No. 9,990,304 ("Samsung's '304 Patent").

419. During the prosecution of Samsung's '304 Patent, Samsung Electronics Co., Ltd. cited Radian's '719 Publication.

420. Radian's '719 Publication issued as Radian's '454 Patent.

421. The '183 Patent asserted by Radian in this action is a descendant of Radian's '454 Patent.

422. On November 13, 2015, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 14/941,517 ("Samsung's '517 Application"), entitled "Selective underlying exposure storage mapping," which issued on June 12, 2018, as U.S. Pat. No. 9,996,473 ("Samsung's '473 Patent").

423. During the prosecution of Samsung's '473 Patent, Samsung Electronics Co., Ltd. cited Radian's '719 Publication.

424. Radian's '719 Publication issued as Radian's '454 Patent.

425. The '183 Patent asserted by Radian in this action is a descendant of Radian's '454 Patent.

426. On November 13, 2015, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 14/941,525 ("Samsung's '525 Application"), entitled "Multimode storage device," which issued on April 10, 2018, as U.S. Pat. No. 9,940,028 ("Samsung's '028 Patent").

427. During prosecution of Samsung's '028 Patent, Samsung Electronics Co., Ltd. cited Radian's '719 Publication.

428. Radian's '719 Publication issued as Radian's '454 Patent.

429. The '183 Patent asserted by Radian in this action is a descendant of Radian's '454 Patent.

430. On December 28, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/392,353 ("Samsung's '353 Application"), entitled "Storage device including nonvolatile

memory device and controller, operating method of storage device, and method for accessing storage device," which issued on March 12, 2019, as U.S. Pat. No. 10,229,051 ("Samsung's '051 Patent").

431. During the prosecution of Samsung's '051 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

432. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

433. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

434. On May 10, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/151,470 ("Samsung's '470 Application"), entitled "User configurable passive background operation," which issued on September 10, 2019, as U.S. Pat. No. 10,409,719 ("Samsung's '719 Patent").

435. During the prosecution of Samsung's '719 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

436. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

437. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

438. On June 17, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/186,250 ("Samsung's '250 Application"), entitled "Mechanism for SSDs to efficiently manage background activity with notify," which issued on October 23, 2018, as U.S. Pat. No. 10,108,450 ("Samsung's '450 Patent").

439. During the prosecution of Samsung's '450 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

440. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

441. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

442. On August 19, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/242,433 ("Samsung's '433 Application"), entitled "Morphic storage device," which issued on August 21, 2018, as U.S. Pat. No. 10,055,159 ("Samsung's '159 Patent").

443. During the prosecution of Samsung's '159 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

444. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

445. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

446. On August 2, 2018, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 16/053,686 ("Samsung's '686 Application"), entitled "Morphic storage device," which issued on November 26, 2019, as U.S. Pat. No. 10,489,076 ("Samsung's '076 Patent").

447. During the prosecution of Samsung's '076 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

448. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

449. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376

Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

450. On January 20, 2017, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/411,962 ("Samsung's '962 Application"), entitled "Multi-mode NVMe over fabrics devices," which issued on August 6, 2019 as U.S. Pat. No. 10,372,659 ("Samsung's '659 Patent").

451. During the prosecution of Samsung's '659 Patent, Samsung Electronics Co., Ltd. cited Radian's '749 Patent.

452. Radian's '749 Patent, which was cited by Samsung Electronics Co., Ltd. is a continuation of Radian's '376 Patent, from which the '183 Patent asserted by Radian is descended.

453. On November 7, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/345,509 ("Samsung's '509 Application"), entitled "Self-configuring baseboard management controller (BMC)," which issued on October 4, 2022, as U.S. Pat. No. 11,461,258 ("Samsung's '258 Patent").

454. During the prosecution of Samsung's '258 Patent, Samsung Electronics Co., Ltd. cited Radian's '749 Patent.

455. Radian's '749 Patent, which was cited by Samsung Electronics Co., Ltd. is a continuation of Radian's '376 Patent, from which the '183 Patent asserted by Radian is descended.

456. On April 23, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 16/857,172 ("Samsung's '172 Application"), entitled "Self-configuring SSD multi-protocol support in host-less environment," which issued on October 12, 2021, as U.S. Pat. No. 11,144,496 ("Samsung's '496 Patent").

457. During the prosecution of Samsung's '496 Patent, Samsung Electronics Co., Ltd. cited Radian's '749 Patent.

458. Radian's '749 Patent, which was cited by Samsung Electronics Co., Ltd. is a

continuation of Radian's '376 Patent, from which the '183 Patent asserted by Radian is descended.

459. On January 10, 2017, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/403,088 ("Samsung's '088 Application"), entitled "System and method for supporting multi-path and/or multi-mode NMVe over fabrics devices," which issued on February 19, 2019, as U.S. Pat. No. 10,210,123 ("Samsung's '123 Patent").

460. During the prosecution of Samsung's '123 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

461. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

462. On November 7, 2016, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/345,507 ("Samsung's '507 Application"), entitled "Method for using BMC as proxy NVMeoF discovery controller to provide NVM subsystems to host," which issued on July 9, 2019 as U.S. Pat. No. 10,346,041 ("Samsung's '041 Patent").

463. During the prosecution of Samsung's '041 Patent, Samsung Electronics Co., Ltd. cited Radian's '749 Patent.

464. Radian's '749 Patent, which was cited by Samsung Electronics Co., Ltd. is a continuation of Radian's '376 Patent, from which the '183 Patent asserted by Radian in this action is descended.

465. On April 6, 2017, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/481,147 ("Samsung's '147 Application"), entitled "Flash-integrated high bandwidth memory appliance," which issued on April 6, 2017, as U.S. Pat. No. 11,397,687 ("Samsung's '687 Patent").

466. During the prosecution of Samsung's '687 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

467. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

468. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

469. On August 25, 2017, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 15/686,690 ("Samsung's '690 Application"), entitled "System and method for managing memory device." The '690 Application was published on December 20, 2018, as U.S. Pat. Pub. No. 2018/0364937, and was ultimately abandoned.

470. During the prosecution of Samsung's '690 Application, the examiner cited Radian's '376 Patent.

471. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '690 Application.

472. On July 10, 2019, Samsung Electronics Co., Ltd. filed German Pat. App. No. 10 2019 118 631.5 ("Samsung's '631.5 Application"), entitled "Storage device that initiates maintenance work independently without command from the host and electronic system with the same." Samsung's '631.5 Application was published on February 6, 2020, as German Pat. Pub. No. 10 2019 118 631 and remains pending.

473. During the prosecution of Samsung's '631.5 Application, the examiner cited Radian's '376 Patent.

474. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '631.5 Application.

475. The examiner also cited Radian's U.S. Pat. No. 9,542,118 ("the '118 Patent").

476. The '656, '657, and '995 Patents asserted by Radian in this action are descendants

of Radian's '118 Patent, which was cited against Samsung's '631.5 Application.

477. The examiner also cited Radian's U.S. Pat. No. 10,552,058 ("the '058 Patent").

478. Radian's '058 Patent issued from U.S. Pat. App. No. 15/211,927, and claims priority from U.S. Prov. Pat. App. No. 62/199,969 and U.S. Prov. Pat. App. No. 62/194,172.

479. The '656, '657, and '995 Patents asserted by Radian in this action also claim priority from U.S. Prov. Pat. App. No. 62/199,969 and U.S. Prov. Pat. App. No. 62/194,172.

480. On June 25, 2020, Samsung Electronics Co. filed European Pat. App. No. 20182197.2 ("Samsung's '197.2 Application"), entitled "Storage device and operating method of storage device," which issued as European Pat. No. 3,771,983 ("Samsung's '983 Patent").

481. During the prosecution of Samsung's '983 Patent, Samsung Electronics Co., Ltd. cited Radian's '854 Patent.

482. Radian's '854 Patent, which was cited by Samsung Electronics Co., Ltd. is a continuation of Radian's '376 Patent.

483. During the prosecution of Samsung's '983 Patent, the examiner cited Radian's '376 Patent.

484. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '197.2 Application.

485. On September 1, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/009,684 ("Samsung's '684 Application"), entitled "Storage device block-level failure prediction-based data placement," which issued on August 22, 2023, as U.S. Pat. No. 11,734,093 ("Samsung's '093 Patent").

486. During the prosecution of Samsung's '093 Patent, the examiner cited Radian's U.S. Pat. 11,354,234 ("the '234 Patent").

487. Radian's '234 Patent, which was cited against Samsung's '684 Application, is a descendant of Radian's '376 Patent.

488. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent.

489. On September 17, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/024,598 ("Samsung's '598 Application"), entitled "Systems and methods for processing commands for storage devices," which issued on August 22, 2023, as U.S. Pat. No. 11,733,918 ("the '918 Patent").

490. During the prosecution of Samsung's '918 Patent, Samsung Electronics Co., Ltd. cited Radian's '129 Publication.

491. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's ' 376 Patent.

492. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited by Samsung Electronics Co., Ltd.

493. On October 2, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/062,467 ("Samsung's '467 Application"), entitled "Systems and methods for processing copy commands," which issued on October 17, 2023, as U.S. Pat. No. 11,789,634 ("Samsung's '634 Patent").

494. During the prosecution of Samsung's '634 Patent, the examiner cited Radian's '129 Publication.

495. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

496. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376

Patent, the published application of which was cited against Samsung's '467 Application.

497. On July 7, 2021, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/369,884 ("Samsung's '884 Application"), entitled "Method of operating memory device and host device, and memory system including partitioning purge region responsive to purge information," which issued on October 24, 2023, as U.S. Pat. No. 11,797,210 ("Samsung's '210 Patent").

498. During the prosecution of Samsung's '210 Patent, the examiner cited Radian's '129 Publication.

499. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

500. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited against Samsung's '884 Application.

501. On May 5, 2021, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/308,991 ("Samsung's '991 Application"), entitled "Systems, methods, and devices for data storage with specified data transfer rate," which issued on August 15, 2023, as U.S. Pat. No. 11,726,659 ("Samsung's '659 Patent").

502. During the prosecution of Samsung's '659 Patent, the examiner cited Radian's '376 Patent.

503. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '991 Application.

504. On August 3, 2021, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/393,399 ("the '399 Application"), entitled "Systems, methods, and apparatus for page migration in memory systems." The '399 Application was published on December 1, 2022 as U.S. Pat. Pub. No. 2022/0382478 and remains pending.

505. During the prosecution of Samsung's '399 Application, the examiner cited Radian's U.S. Pat. No. 10,642,505 ("the '505 Patent").

506. Radian's '505 Patent, which was cited against Samsung's '399 Application, is a descendant of Radian's '376 Patent.

507. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent.

508. During the prosecution of Samsung's '399 Application, the examiner also cited U.S. Pat. No. 11,347,639 ("the '639 Patent"), another Radian patent.

509. Radian's '639 Patent, which was cited against Samsung's '399 Application, is a descendant of Radian's '376 Patent.

510. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent.

511. During the prosecution of Samsung's '399 Application, the examiner also cited Radian's U.S. Pat. No. 11,275,695 ("the '695 Patent"),

512. Radian's '695 Patent, which was cited against Samsung's '399 Application, is a descendant of Radian's U.S. Pat. No. 10,642,748 ("the '748 Patent").

513. The '656, '657, and '995 Patents asserted by Radian in this action are also descendants of Radian's '748 Patent.

514. On January 19, 2023, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 18/099,246 ("Samsung's '246 Application"), entitled "Systems, methods, and apparatus for data placement in a storage device." Samsung's '246 Application was published on December 1, 2022, as U.S. Pat. Pub. No. 2024/0012579 and remains pending.

515. During the prosecution of Samsung's '246 Application, the examiner cited

Radian's '719 Publication.

516. Radian's '719 Publication, which was cited against Samsung's '246 Application, was the published application of Radian's U.S. Pat. App. No. 14/466,167, which issued as Radian's '454 Patent.

517. The '183 Patent asserted by Radian in this action is a descendant of Radian's '454 Patent.

518. On November 28, 2013, Samsung Electronics Co., Ltd. filed Korea Pat. App. No. 10-2013-0146321 ("Samsung's '321 Application"), entitled "All-in-one data storage device having internal hardware filter, method thereof, and system having the data storage device," which issued on May 29, 2015, as Korean Pat. App. No. 10-2103543 ("Samsung's '543 Patent").

519. During the prosecution of Samsung's '543 Patent, the examiner cited Radian's '376 Patent.

520. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '321 Application.

521. On January 10, 2014, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2014-0003245 ("Samsung's '245 Application"), entitled "Method for processing data on storage device and storage device," which issued on November 23, 2020, as Korean Pat. No. 10-2181210 ("the '210 Patent").

522. During the prosecution of Samsung's '210 Patent, the examiner cited Radian's '376 Patent.

523. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '245 Application.

524. On August 18, 2018, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-

2014-0107128 ("Samsung's '128 Application"), entitled "Operating method of memory controller and nonvolatile memory system including the memory controller," which issued on August 31, 2020, as Korean Pat. No. 10-2148889 ("Samsung's '889 Patent").

525. During the prosecution of Samsung's '889 Patent, the examiner cited Radian's '376 and '854 Patents.

526. The '854 Patent cited against Samsung's '128 Application is a continuation of Radian's '376 Patent.

527. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '128 Application.

528. On June 8, 2015, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2015-0080741 ("Samsung's '741 Application"), entitled "Nonvolatile memory module, computing system having the same, and PVT compensation method thereof," which issued on November 2, 2022, as Korean Pat. No. 10-2461460 ("Samsung's '460 Patent").

529. During the prosecution of Samsung's '460 Patent, the examiner cited Radian's '129 Publication.

530. Radian's '129 Publication was the published application of Radian's '723 Application, which later issued as Radian's '376 Patent.

531. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, the published application of which was cited against Samsung's '741 Application.

532. On April 1, 2016, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2016-0040320 ("Samsung's '320 Application"), entitled "Storage device and event notification method thereof," which issued on June 30, 2023, as Korean Pat. No. 10-2549611 ("Samsung's '611 Patent").

533. During the prosecution of Samsung's '611 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

534. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

535. On August 4, 2016, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2016-0099599 ("Samsung's '599 Application"), entitled "Storage device using host memory and operating method thereof," which issued on May 23, 2024, as Korean Pat. No. 10-2667430 ("Samsung's '430 Patent").

536. During the prosecution of Samsung's '430 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

537. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

538. On November 24, 2017, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2017-0158832 ("Samsung's '832 Application"), entitled "Method of managing data and storage device performing the same," which issued on August 16, 2023, as Korean Pat. No. 10-2567140 ("Samsung's '140 Patent").

539. During the prosecution of Samsung's '140 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

540. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

541. During the prosecution of Samsung's '140 Patent, Samsung Electronics Co., Ltd. also cited Radian's '118 Patent.

542. The '656, '657, and '995 Patents asserted by Radian in this action are descendants

of Radian's '118 Patent, which was cited by Samsung Electronics Co., Ltd.

543. On July 24, 2018, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2018-0085788 ("Samsung's '788 Application"), entitled "Solid state drive and a method for metadata access," which issued on June 28, 2023, as Korean Pat. No. 10-2549346 ("Samsung's '346 Patent").

544. During the prosecution of Samsung's '346 Patent, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

545. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

546. On March 20, 2019, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2019-0031649 ("Samsung's '649 Application"), entitled "Operation method of open-channel storage device." Samsung's '649 Application was published on October 6, 2020, as Korean Pat. Pub. 10-2020-0113047.

547. During the prosecution of Samsung's '649 Application, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

548. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

549. On October 2, 2019, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2019-0122655 ("Samsung's '655 Application"), entitled "Host System managing assignment of free block, Data Processing System having the same and Operating Method of Host System." Samsung's '655 Application was published on April 12, 2021, as Korean Pat. Pub. No. 10-2021-0039872, and remains pending.

550. During the prosecution of Samsung's '655 Application, the examiner cited

Radian's '376 Patent.

551. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '655 Application.

552. On October 4, 2019, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2019-0123347 ("Samsung's '347 Application"), entitled "Operating method of memory system and host recovering data with write error." Samsung's '347 Application was published on April 15, 2021, as Korean Pat. Pub. No. 10-2021-0041158, and remains pending.

553. During the prosecution of Samsung's '347 Application, Samsung Electronics Co., Ltd. cited Radian's '854 Patent.

554. The '854 Patent, which was cited by Samsung Electronics Co., Ltd., is a continuation of Radian's '376 Patent.

555. During the prosecution of Samsung's '347 Application, the examiner cited Radian's '376 Patent.

556. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '347 Application.

557. On October 4, 2019, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2019-0123348 ("Samsung's '348 Application"), entitled "Operating method of memory system and host recovering data with correctable read error." Samsung's '348 Application was published on April 15, 2021, as Korean Pat. Pub. No. 10-2021-0041159, and remains pending.

558. During the prosecution of Samsung's '348 Application, Samsung Electronics Co., Ltd. cited Radian's '376 Patent.

559. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited by Samsung Electronics Co., Ltd.

560. On December 2, 2019, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2019-0158001 ("Samsung's '001 Application"), entitled "Storage device, storage system and method of operating storage device." Samsung's '001 Application was published on June 10, 2021, as Korean Pat. Pub. No. 10-2021-0068699, and remains pending.

561. During the prosecution of Samsung's '001 Application, the examiner cited Radian's '376 Patent.

562. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '001 Application.

563. During the prosecution of Samsung's '001 Application, the examiner cited Radian's '058 Patent.

564. Radian's '058 Patent, which was cited against Samsung's '001 Application, claims priority from two of the same provisional applications from which the '656, '657, and '995 Patents asserted by Radian in this action claim priority.

565. On November 16, 2020, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2020-0152806 ("Samsung's '806 Application"), entitled "Storage device, electronic system including the same storage device and the method of operating the same storage device." Samsung's '806 Application was published on May 24, 2022, as Korean Pat. Pub. No. 10-2022-0066601, and remains pending.

566. During the prosecution of Samsung's '806 Application, the examiner cited Radian's '376 Patent.

567. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '806 Application.

568. On November 23, 2020, Samsung Electronics Co., Ltd. filed Korean Pat. App. No.

10-2020-0158053 ("Samsung's '053 Application"), entitled "Memory device, system including the same and operating method of memory device." Samsung's '053 Application was published on May 31, 2022, as Korean Pat. Pub. No. 10-2022-0070951, and remains pending.

569. During the prosecution of Samsung's '053 Application, the examiner cited Radian's '376 Patent.

570. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '053 Application.

571. On March 2, 2021, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2021-0027516 ("Samsung's '516 Application"), entitled "Storage controller redirecting a write operation and operating method thereof." Samsung's '516 Application was published on September 14, 2022, as Korean Pat. Pub. No. 10-2022-0124318, and remains pending.

572. During the prosecution of Samsung's '516 Application, the examiner cited Radian's '376 Patent.

573. The '183 Patent asserted by Radian in this action is a descendant of Radian's '376 Patent, which was cited against Samsung's '516 Application.

574. On July 2, 2018, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 16/026,008 ("Samsung's '008 Application"), entitled "Apparatus to insert error-correcting coding (ECC) information as data within dynamic random access memory (DRAM)," which issued on December 1, 2020 as U.S. Pat. No. 10,853,168 ("Samsung's '168 Patent").

575. During the prosecution of Samsung's '516 Application, the examiner cited Radian's U.S. Pat. No. 10,445,229 ("the '229 Patent").

576. The '772 and '614 Patents asserted by Radian in this action are descendants of Radian's '229 Patent, which was cited against Samsung's '008 Application.

577. On February 3, 2020, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2020-0012629 ("Samsung's '629 Application"), entitled "Stacked memory device, and operating method thereof." Samsung's '629 Application was published on August 11, 2021, as Korean Pat. Pub. No. 10-2021-0098728, and remains pending.

578. During the prosecution of Samsung's '516 Application, the examiner cited Radian's U.S. Pat. No. 10,445,229 ("the '229 Patent").

579. The '772 and '614 Patents asserted by Radian in this action are descendants of Radian's '229 Patent, which was cited against Samsung's '516 Application.

580. On June 19, 2018, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 16/012,470 ("Samsung's '470 Application"), entitled "System and method for optimizing performance of a solid-state drive using a deep neural network," which issued on March 30, 2021, as U.S. Pat. No. 10,963,394 ("Samsung's '394 Patent").

581. During the prosecution of Samsung's '394 Patent, Samsung Electronics Co., Ltd. cited Radian's U.S. Pat. No. 9,785,572 ("the '572 Patent").

582. The '656, '657, and '995 Patents asserted by Radian in this action are descendants of Radian's '572 Patent, which was cited by Samsung Electronics Co., Ltd.

583. On October 22, 2020, Samsung Electronics Co., Ltd. filed U.S. Pat. App. No. 17/077,200 ("Samsung's '200 Application"), entitled "Operating method of memory system including memory controller and nonvolatile memory device," which issued on April 19, 2022, as U.S. Pat. No. 11,309,032 ("Samsung's '032 Patent").

584. During the prosecution of Samsung's '032 Patent, the examiner cited Radian's '118 Patent.

585. The '656, '657, and '995 Patents asserted by Radian in this action are descendants

of the ’118 Patent, which was cited against Samsung’s ’200 Application.

586. During the prosecution of Samsung’s ’032 Patent, the examiner also cited Radian’s ’058 Patent.

587. Radian’s ’058 Patent, which was cited against Samsung’s ’200 Application, claims priority from U.S. Prov. Pat. App. No. 62/199,969 and U.S. Prov. Pat. App. No. 62/194,172; the ’656, ’657, and ’995 Patents asserted by Radian in this action also claim priority to U.S. Prov. Pat. App. No. 62/199,969 and U.S. Prov. Pat. App. No. 62/194,172.

588. On August 4, 2016, Samsung Electronics Co., Ltd. filed Korean Pat. App. No. 10-2016-0099219 (“Samsung’s ’219 Application”), entitled “Nonvolatile memory device,” which issued on January 3, 2024, as Korean Pat. No. 10-2620562 (“Samsung’s ’562 Patent”).

589. During the prosecution of Samsung’s ’562 Patent, the examiner cited Radian’s ’118 Patent.

590. The ’656, ’657, and ’995 Patents asserted by Radian in this action are descendants of Radian’s ’118 Patent, which was cited against Samsung’s ’219 Application.

J. Samsung Entity Relationships

591. SEC is the flagship entity of the Samsung Group, a South Korean *chaebol* comprised of approximately 63 affiliated companies as of 2023.[19]

592. SEC is the world’s largest manufacturer of semiconductor memory, and it is also

[19] WENDOVER PRODUCTIONS, *Samsung’s Dangerous Dominance over South Korea*, available at https://www.youtube.com/watch?v=oL0umpPPe-8 at 04:48–04:59 (“These family-controlled, government-aided, corporate groups came to be known as chaebols—a term that entered common Korean parlance during the ‘60s and ‘70s as the institutions grew.”); *see also* Statista, *Number of affiliates of leading business groups in South Korea as of May 2023, by business group*, available at https://www.statista.com/statistics/1314469/south-korea-number-of-subsidiaries-of-conglomerates-by-company/.

one of the world's leading manufacturers of consumer electronics, electronic components, and semiconductors.

593. SEA is one of SEC's wholly owned subsidiaries.

594. SEA exists for the purpose of commercializing, marketing, selling, distributing, and servicing electronics that are developed, designed, and manufactured by SEC for sale in the United States.

595. SEA is involved in commercializing, marketing, selling, distributing, and servicing consumer electronics developed, designed, and manufactured by SEC.

596. SEA does not itself manufacture any products.

597. SEA markets consumer electronics designed for use in the U.S. to customers and end-users in the U.S.

598. SEC is responsible for designing the products sold by SEA specifically for use in the United States.

599. SEC manufactures these products that it has specifically designed for use in the United States and provides them to SEA for the express purpose of being sold in the United States by SEC.

600. SEC and SEA have entered into agreements regarding the distribution, sales, and service of SEC-manufactured products by SEA in the United States.

601. SEC and SEA did not enter into these agreements at arms' length as equals.

602. The terms of these agreements were dictated by SEC to SEA.

603. The distribution, sales, and service of SEC-manufactured products in the United States by SEA pursuant to these agreements therefore necessarily occur at SEC's direction.

604. SEC exercises effective control over SEA's activities.

605. SEA acts merely as, and exists for the purpose of acting as, an agent through which SEC conducts business in the United States.

606. The degree of control exercised by SEC over SEA is greater than that normally associated with common ownership and directorship.

607. This degree of control extends to SEA's own subsidiaries.

608. SEC exercises effective control over SEA's subsidiaries.

609. SEC does not allow SEA to take licenses that involve payment for unpaid royalties without SEC's approval.

610. This is because such payments might be deemed by the South Korean government as an attempt to circumvent the tax it levies against payments for licensing U.S. patents.

611. When SEC agrees to pay United States companies for unpaid royalties (whether for its own infringement or for the infringement of one of its subsidiaries), SEC withholds South Korean taxes on such payments.

612. Employees and officers of any given Samsung entity act on behalf of other Samsung entities, all of whom are ultimately under SEC's control and direction.

613. Samsung employees and officers of one Samsung entity attend industry conferences on behalf of or while being identified as being associated with other Samsung entities.

614. Samsung employees and officers of one Samsung entity present papers, lectures, or presentations on behalf of or while being identified as being associated with other Samsung entities.

615. Samsung employees and officers of one Samsung entity attend industry organizations on behalf of or while being identified as being associated with other Samsung entities.

616. Samsung employees and officers of one Samsung entity post articles, papers, press releases, blog posts, videos, and other content on behalf of or while being identified as being associated with other Samsung entities.

617. Samsung employees and officers of one Samsung entity post articles, papers, press releases, blog posts, videos, and other content on the websites of other Samsung entities.

618. Samsung entities post product information, advertising, marketing, promotional information, data sheets, specification sheets, manuals, instructions, guides, and other literature on each other's websites.

619. Samsung Electronics Co., Ltd. controls and operates the "samsung.com" domain from its headquarters in the Republic of Korea.[20]

620. The "samsung.com" domain from its headquarters in the Republic of Korea is registered to Samsung Electronics Co., Ltd. with an address in the Republic of Korea.[21]

621. Samsung Electronics Co., Ltd. controls and operates other Samsung subsidiaries' websites from its headquarters in the Republic of Korea.[22]

K. Samsung Electronics Co., Ltd.'s Vicarious Liability

622. SEC employees and officers are agents of SEC.

623. SEC directs and controls the actions and performance of SEC employees and

---

[20] *See* WHOIS.COM, *Domain Information, Samsung.com*, available at https://www.whois.com/whois/samsung.com (identifying Samsung Electronics Co., Ltd., with an address in Korea, as the registrant of the domain "samsung.com").
[21] *Id.*
[22] *E.g.*, https://semiconductor.samsung.com/us/legal/ ("Samsung controls and operates this site from it's (sic) headquarters in the Republic of Korea."); *see also* https://www.samsungcnt.com/eng/legal-notice.do ("Any disputes arising between the user and Samsung C&T Corporation in regard to use of this website shall resort to Seoul Central District Court or a competent civil court.").

officers in the course of their employment, including those related to infringement of the Asserted Patents.

624. SEC authorizes the actions and performance of SEC employees and officers, including those related to infringement of the Asserted Patents.

625. SEC approves the actions and performance of SEC employees and officers, including those related to infringement of the Asserted Patents.

626. SEC specifies the timing and manner of the performance of activities by SEC employees and officers, including those related to infringement of the Asserted Patents.

627. SEC profits from the activities of SEC employees and officers.

628. SEC has the rights, powers, or abilities to cause SEC employees and officers to stop or limit their infringing activities.

629. SEC has not exercised its rights, powers, or abilities to cause SEC employees and officers to stop or limit their infringing activities.

630. SEC is vicariously liable for the infringing activities of SEC employees and officers.

631. SEA is an agent of SEC.

632. SEC directs and controls the actions and performance of SEA, including those related to infringement of the Asserted Patents.

633. SEC authorizes the actions and performance of SEA, including those related to infringement of the Asserted Patents.

634. SEC approves the actions and performance of SEA, including those related to infringement of the Asserted Patents.

635. SEC conditions benefits derived by SEA on the performance of activities, including

those related to infringement of the Asserted Patents.

636. SEC specifies the timing and manner of the performance of activities by SEA, including those related to infringement of the Asserted Patents.

637. SEC profits from the activities of SEA.

638. SEC has the rights, powers, or abilities to cause SEA to stop or limit its infringing activities.

639. SEC has not exercised its rights, powers, or abilities to cause SEA to stop or limit its infringing activities.

640. SEC is vicariously liable for the infringing activities of SEA.

641. SEA employees and officers are agents of SEC.

642. SEC directs and controls the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

643. SEC authorizes the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

644. SEC approves the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

645. SEC specifies the timing and manner of the performance of activities by SEA employees and officers, including those related to infringement of the Asserted Patents.

646. SEC profits from the activities of SEA employees and officers.

647. SEC has the rights, powers, or abilities to cause SEA employees and officers to stop or limit their infringing activities.

648. SEC has not exercised its rights, powers, or abilities to cause SEA employees and officers to stop or limit their infringing activities.

649. SEC is vicariously liable for the infringing activities of SEA employees and officers.

650. Samsung subsidiaries other than SEA ("other Samsung subsidiaries") are agents of SEC.

651. SEC directs and controls the actions and performance of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

652. SEC authorizes the actions and performance of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

653. SEC approves the actions and performance of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

654. SEC conditions benefits derived by other Samsung subsidiaries on the performance of activities, including those related to infringement of the Asserted Patents.

655. SEC specifies the timing and manner of the performance of activities by other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

656. SEC profits from the activities of other Samsung subsidiaries.

657. SEC has the rights, powers, or abilities to cause other Samsung subsidiaries to stop or limit its infringing activities.

658. SEC has not exercised its rights, powers, or abilities to cause other Samsung subsidiaries to stop or limit its infringing activities.

659. SEC is vicariously liable for the infringing activities of other Samsung subsidiaries.

660. The employees and officers of other Samsung subsidiaries are agents of SEC.

661. SEC directs and controls the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

662. SEC authorizes the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

663. SEC approves the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

664. SEC specifies the timing and manner of the performance of activities by the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

665. SEC profits from the activities of the employees and officers of other Samsung subsidiaries.

666. SEC has the rights, powers, or abilities to cause the employees and officers of other Samsung subsidiaries to stop or limit their infringing activities.

667. SEC has not exercised its rights, powers, or abilities to cause the employees and officers of other Samsung subsidiaries to stop or limit their infringing activities.

668. SEC is vicariously liable for the infringing activities of the employees and officers of other Samsung subsidiaries.

669. SEC conditions benefits derived by Samsung customers on the performance of activities, including those related to infringement of the Asserted Patents.

670. SEC specifies the timing and manner of the performance of activities by Samsung customers, including those related to infringement of the Asserted Patents.

671. SEC profits from the activities of Samsung customers.

672. SEC has the rights, powers, or abilities to cause Samsung customers to stop or limit its infringing activities.

673. SEC has not exercised its rights, powers, or abilities to cause Samsung customers

to stop or limit its infringing activities.

674. SEC is vicariously liable for the infringing activities of Samsung customers.

L. Samsung Electronics America, Inc.'s Vicarious Liability

675. SEA employees and officers are agents of SEA.

676. SEA directs and controls the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

677. SEC authorizes the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

678. SEC approves the actions and performance of SEA employees and officers, including those related to infringement of the Asserted Patents.

679. SEA specifies the timing and manner of the performance of activities by SEA employees and officers, including those related to infringement of the Asserted Patents.

680. SEA profits from the activities of SEA employees and officers.

681. SEA has the rights, powers, or abilities to cause SEA employees and officers to stop or limit their infringing activities.

682. SEA has not exercised its rights, powers, or abilities to cause SEA employees and officers to stop or limit their infringing activities.

683. SEA is vicariously liable for the infringing activities of SEA employees and officers.

684. Other Samsung subsidiaries are agents of SEA.

685. SEA directs and controls the actions and performance of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

686. SEA authorizes the actions and performance of other Samsung subsidiaries,

including those related to infringement of the Asserted Patents.

687. SEA approves the actions and performance of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

688. SEA conditions benefits derived by other Samsung subsidiaries on the performance of activities, including those related to infringement of the Asserted Patents.

689. SEA specifies the timing and manner of the performance of activities by other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

690. SEA profits from the activities of other Samsung subsidiaries.

691. SEA has the rights, powers, or abilities to cause other Samsung subsidiaries to stop or limit its infringing activities.

692. SEA has not exercised its rights, powers, or abilities to cause other Samsung subsidiaries to stop or limit its infringing activities.

693. SEA is vicariously liable for the infringing activities of other Samsung subsidiaries.

694. SEA authorizes the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

695. SEA approves the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

696. The employees and officers of other Samsung subsidiaries are agents of SEA.

697. SEA directs and controls the actions and performance of the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

698. SEA specifies the timing and manner of the performance of activities by the employees and officers of other Samsung subsidiaries, including those related to infringement of the Asserted Patents.

699. SEA profits from the activities of the employees and officers of other Samsung subsidiaries.

700. SEA has the rights, powers, or abilities to cause the employees and officers of other Samsung subsidiaries to stop or limit their infringing activities.

701. SEA has not exercised its rights, powers, or abilities to cause the employees and officers of other Samsung subsidiaries to stop or limit their infringing activities.

702. SEA is vicariously liable for the infringing activities of the employees and officers of other Samsung subsidiaries.

703. SEA conditions benefits derived by Samsung customers on the performance of activities, including those related to infringement of the Asserted Patents.

704. SEA specifies the timing and manner of the performance of activities by Samsung customers, including those related to infringement of the Asserted Patents.

705. SEA profits from the activities of Samsung customers.

706. SEA has the rights, powers, or abilities to cause Samsung customers to stop or limit its infringing activities.

707. SEA has not exercised its rights, powers, or abilities to cause Samsung customers to stop or limit its infringing activities.

708. SEA is vicariously liable for the infringing activities of Samsung customers.

JURISDICTION AND VENUE

709. Radian incorporates the paragraphs of the previous sections into the following section by reference.

710. Radian's claims for patent infringement against Defendants arise under the Patent Act, 35 U.S.C. § 1 et seq.

711. Subject matter jurisdiction is proper in this Court under 28 U.S.C. §§ 1331 and 1338(a).

712. SEC and SEA are subject to the personal jurisdiction of this Court.

713. SEC and SEA have not disputed this Court's personal jurisdiction over them in recent patent infringement actions filed in this District. *See, e.g.*, *Mobile Data Technologies LLC v. Samsung Elecs. Am., Inc. et al.*, No. 2:24-cv-00435, Dkt. 26 at ¶ 7 (E.D. Tex. Oct. 7, 2024); *Four Batons Wireless, LLC v. Samsung Elecs. Co., Ltd. et al.*, No. 2:24-cv-00284, Dkt. 24 at ¶ 8 (E.D. Tex. Aug. 19, 2024); *Cerrence Operating Co. v. Samsung Elec. Co., Ltd. et al.*, No. 2:24-cv-00181-JRG, Dkt. 25 at ¶ 8 (E.D. Tex. July 10, 2024).

714. SEC and SEA have sufficient minimum contacts and/or have engaged in continuous and systematic activities in the forum as a result of business conducted within the State of Texas and the Eastern District of Texas.

715. SEA, directly or through subsidiaries, makes, uses, sells, offers for sale, imports, advertises, makes available, and/or markets products within the State of Texas and the Eastern District of Texas that infringe one or more claims of the Asserted Patents.

716. SEA has done these activities at the direction of its owner, SEC.

717. SEA has placed or contributed to placing infringing products into the stream of commerce knowing or understanding that such products would be sold and used in the United States, including in this District.

718. SEA has done so at the direction of its owner, SEC, which also knew or understood that such products would be sold and used in the United States, including in this District.

719. SEA has derived substantial revenues from infringing acts in the Eastern District of Texas, including from the sale and use of infringing products.

720. These revenues have substantially flowed to SEC, as the sole owner of SEA.

721. Venue in this District is proper under 28 U.S.C. § 1391(c)(3) and 28 U.S.C. § 1400(b).

722. SEC is not a resident of the United States and may be sued in this District because suits against foreign entities are proper in any judicial district where they are subject to personal jurisdiction. 28 U.S.C. § 1391(c)(3); *In re HTC Corp.*, 889 F.3d 1349, 1356–61 (Fed. Cir. 2018).

723. SEA maintains a regular and established place of business in this judicial district at 6625 Excellence Way, Plano, Texas 75023.

724. SEC and SEA have admitted in previous patent infringement actions filed against them that venue in this District is proper. *See, e.g.*, *Four Batons*, Dkt. 24 at ¶ 15; *Cerence*, Dkt. 25 at ¶ 9.

725. SEA has committed acts of infringement in this district.

726. SEA employees in Plano are responsible for the marketing and sale of products to customers across the country.

727. In addition to the location at 6625 Excellence Way, SEA has operations at several other locations within this District.

728. The Collin County Central Appraisal District's website shows that SEA maintains business personal property at a number of locations in the county:[23]

[23] From searching the term "Samsung" under the "Owner Name" field at https://www.collincad.org/propertysearch.

	Property ID ↓ Geographic ID ↓	Owner Name	Property Address	Legal Description	2024 Market Value
1	2546587 P-9000-204-3119-1	SAMSUNG ELECTRONICS AMERICA INC	6625 Excellence Way Plano, TX 75023	BPP at 6625 Excellence Way	$16,285,213
2	2539782 P-9000-204-8501-1	SAMSUNG ELECTRONICS AMERICA INC	1100 Klein Rd #00100 Plano, TX 75074	BPP at 1100 Klein Rd	$867,561
3	2662736 P-9000-210-1325-1	SAMSUNG SDS GLOBAL SCL AMERICA INC	6625 Excellence Way Plano, TX 75023	BPP at 6625 Excellence Way	$247,297
4	2702809 P-9000-214-3749-1	SAMSUNG ELECTRONICS AMERICA INC	1751 N Central Expy #0000C McKinney, TX 75070	BPP at 1751 N Central Expy	$13,915
5	2702808 P-9000-214-4019-1	SAMSUNG ELECTRONICS AMERICA INC	2800 Central Expy Plano, TX 75074	BPP at 2800 Central Expy	$9,062
6	2702811 P-9000-214-4020-1	SAMSUNG ELECTRONICS AMERICA INC	190 E Stacy Rd #03000 Allen, TX 75002	BPP at 190 E Stacy Rd	$15,053
7	2702813 P-9000-214-4021-1	SAMSUNG ELECTRONICS AMERICA INC	3333 Preston Rd #00200 Frisco, TX 75034	BPP at 3333 Preston Rd	$10,131
8	2735537 P-9000-216-1494-1	SAMSUNG RESEARCH AMERICA INC	6105 Tennyson Pkwy #00300 Plano, TX 75024	BPP at 6105 Tennyson Pkwy	$9,366,195
9	2833263 P-9000-221-2310-1	SAMSUNG SDS GLOBAL SCL AMERICA INC	3033 W President George Bush Hwy #00250 Plano, TX 75075	BPP at 3033 W President George Bush Hwy	$249,743
10	2835320 P-9000-221-5144-1	SAMSUNG ELECTRONICS AMERICA INC	1005 Placid Ave #00120 Plano, TX 75074	BPP at 1005 Placid Ave	$1,238,753
11	2838537 P-9000-222-0675-1	SAMSUNG ELECTRONICS AMERICA INC	3580 Preston Rd #00100 Frisco, TX 75034	BPP at 3580 Preston Rd	$209,931
12	2850800 P-9000-222-3448-1	SAMSUNG ELECTRONICS AMERICA INC	6555 Excellence Way #00100 Plano, TX 75023	BPP at 6555 Excellence Way	$4,994,718
13	2872831 P-9000-223-2110-1	SAMSUNG ELECTRONICS AMERICA INC	2601 Preston Rd #01214 Frisco, TX 75034	BPP at Stonebriar Mall	$2,309,244

729. SEA also maintains property in Denton County per the Denton County Central Appraisal District's website:[24]

1 - 8 of 8 items

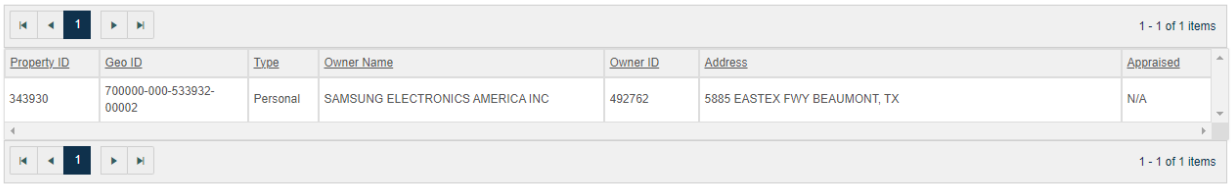
Property ID	Geo ID	Type	Owner Name	Owner ID	Address	Appraised
1020074		Personal	SAMSUNG ELECTRONICS AMERICA INC	1830312	701 LAKESIDE PKWY FLOWER MOUND, TX	$426,641
647771		Personal	SAMSUNG ELECTRONICS AMERICA INC	875087	5299 ELDORADO PKWY FRISCO, TX	$16,776
668025		Personal	SAMSUNG ELECTRONICS AMERICA, INC	904756	6060 LONG PRAIRIE RD 500 FLOWER MOUND, TX 75028	$14,749
668533		Personal	SAMSUNG ELECTRONICS AMERICA INC	906112	2601 S STEMMONS FWY 300 LEWISVILLE, TX	$14,749
668571		Personal	SAMSUNG ELECTRONICS AMERICA, INC	906331	1800 S LOOP 288 380 DENTON, TX 76205	$14,761
680708		Personal	SAMSUNG HVAC AMERICA, LLC	923590	776 HENRIETTA CREEK RD ROANOKE, TX	$5,789,637
731430		Personal	SAMSUNG HVAC AMERICA	992243	776 HENRIETTA CREEK RD 100 ROANOKE, TX 76262-6398	$49,905,943
776168		Personal	SAMSUNG ELECTRONICS AMERICA INC	1057667	800 HENRIETTA CREEK ROANOKE, TX	$37,605,980

1 - 8 of 8 items

730. SEA also maintains property in Jefferson County per the Jefferson County Central

[24] From searching the term "Samsung" under the "Owner Name" field at https://esearch.dentoncad.com/.

Appraisal District's website:[25]

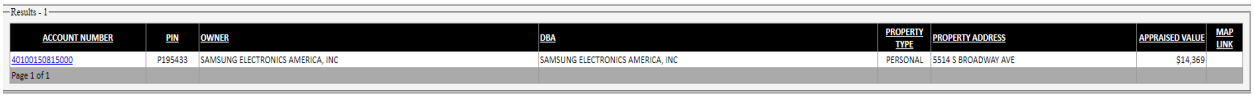


1 - 1 of 1 items

Property ID	Geo ID	Type	Owner Name	Owner ID	Address	Appraised
343930	700000-000-533932-00002	Personal	SAMSUNG ELECTRONICS AMERICA INC	492762	5885 EASTEX FWY BEAUMONT, TX	N/A

1 - 1 of 1 items

731. SEA also maintains property in Smith County per the Smith County Central Appraisal District's website:[26]



Results - 1

ACCOUNT NUMBER	PIN	OWNER	DBA	PROPERTY TYPE	PROPERTY ADDRESS	APPRAISED VALUE	MAP LINK
40100150815000	P195433	SAMSUNG ELECTRONICS AMERICA, INC	SAMSUNG ELECTRONICS AMERICA, INC	PERSONAL	5514 S BROADWAY AVE	$14,369	

Page 1 of 1

732. SEA also maintains property in Grayson County per the Grayson County Central Appraisal District's website:[27]



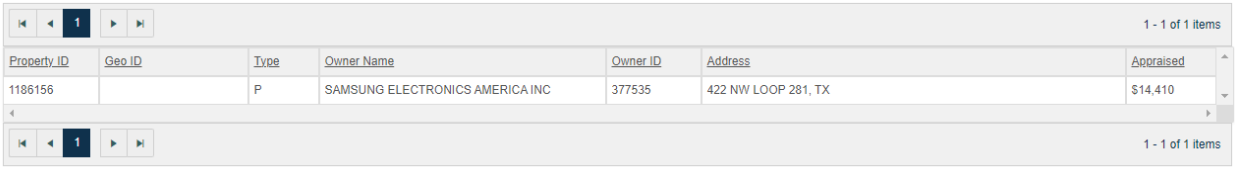
1 - 1 of 1 items

Property ID	Geo ID	Type	Owner Name	Owner ID	Address	Appraised
370473	6P370473	P	SAMSUNG ELECTRONICS AMERICA INC	423640	823 NORTHCREEK DR, TX 75092	N/A

1 - 1 of 1 items

733. SEA also maintains property in Gregg County per the Gregg County Central Appraisal District's website:[28]

[25] From searching the term "Samsung" under the "Owner Name" field at https://esearch.jcad.org/.
[26] From searching the term "Samsung" under the "Owner Name" field at https://www.smithcad.org/Search/PropertySearch.html.
[27] From searching the term "Samsung" under the "Owner Name" field at https://esearch.graysonappraisal.org/.
[28] From searching the term "Samsung" under the "Owner Name" field at https://esearch.gcad.org/.



1 - 1 of 1 items

Property ID	Geo ID	Type	Owner Name	Owner ID	Address	Appraised
1186156		P	SAMSUNG ELECTRONICS AMERICA INC	377535	422 NW LOOP 281, TX	$14,410

1 - 1 of 1 items

734. SEC directs and controls the actions of SEA such that it also maintains regular and established offices in the Eastern District of Texas, including at 6625 Excellence Way, Plano, Texas 75023.

735. SEC directs and controls the actions of SEA personnel commercializing, marketing, selling, distributing, and servicing consumer electronics developed, designed, and manufactured by SEC for use in the United States by customers and end-users in the United States, including SEA personnel located in this District.

COUNT ONE: INFRINGEMENT OF THE '183 PATENT

736. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

737. U.S. Patent No. 11,544,183 ("the '183 Patent"), entitled "Nonvolatile Memory Controller Host-Issued Address Delimited Erasure and Memory Controller Remapping of Host-Address Space for Bad Blocks," was legally and duly issued on January 3, 2023, naming Andrey V. Kuzmin, Mike Jadon, and Richard M. Mathews as the inventor. *See* Exhibit A.

738. Radian owns all rights, title, and interest in the '183 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future infringement.

739. The '183 Patent is valid, enforceable, and directed to patentable subject matter.

740. Radian has complied with 35 U.S.C. § 287 with respect to the '183 Patent.

741. The '183 Patent described challenges faced by conventional memory devices,

including the substantial overhead and write amplification related to flash memory operations. See, for example:

> Such functions typically include caching of write data to reduce frequency of programming operations, wear leveling, bad block management and space reclamation. These tasks are typically managed by a flash memory controller using a flash translation layer (FTL), which keeps records of logical-to-physical translations, wear count, bad blocks and so forth using RAM that is built-in to the flash memory controller.
>
> Each of these functions contributes substantial overhead and write amplification in flash memory. That is to say, substantial data and control bandwidth is consumed in implementing these functions, which can both increase the number of writes to memory (i.e., increase wear) as well as compete with new writes initiated by a host.
>
> What is needed is a mechanism for improving control and data bandwidth for flash memory and other forms of nonvolatile memory. More particularly, a mechanism is needed that reduces control and data bandwidth encumbrances created by memory management functions and thereby decreases the issues referenced above. Still further, a need exists for a memory management scheme that does not create excessive write amplification and bandwidth competition. Finally, a need exists for a flash/nonvolatile memory architecture that has more consistent latency, is conductive to structured pipelining of commands, and permits ubiquitous management of SSDs and other forms of memory in direct-attached and network storage applications.[29]

742. The ’183 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> This disclosure provides techniques for cooperative interaction between a memory controller and host. The techniques call for the memory controller to store information specific to each

[29] Ex. A, U.S. Patent No. 11,544,183, at 1:30–3:37.

> of plural subdivisions of memory, and to make data based on that stored information accessible to the host to assist with management of the memory.[30]
>
> Note that in one embodiment, this infrastructure can be employed to substantially eliminate the need for a flash memory controller to implement a flash translation layer (FTL). That is, a flash memory controller can maintain per-subdivision data, which is accessible to the host (e.g., retrievable by, or against which the host can issue function calls or queried). The host uses this information to issue new write commands that are targeted to specific physical locations in flash memory, thus substantially avoiding the need for translation at a memory controller, and reducing the likelihood of uneven wear.[31]
>
> By redefining host and/or controller responsibilities, host-controller management features duplication and associated disadvantages can also be avoided, leading to a simpler and less expensive memory controller design.[32]

743. Some of these described solutions are reflected in the claimed inventions of the ’183 Patent.

744. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the ’183 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the ’183 Patent.

745. Claim 1 of the ’183 Patent recites:

> [preamble] A memory controller to control flash memory, the flash memory having

---

[30] *Id*. at 4:46–51.

[31] *Id*. at 5:27–37.

[32] *Id*. at 6:33–37.

subdivisions, each of the subdivisions corresponding to a respective group of erase blocks, each of the erase blocks such that it must be erased before storage locations in that erase block can be reprogrammed, the memory controller comprising:

[1a] a host interface;

[1b] a memory interface;

[1c] logic to track information indicating extent of page utilization of each of the subdivisions;

[1d] logic to receive a write request from a host via the host interface, the write request accompanied by an address that designates a specific one of the subdivisions, wherein the logic to track is to, in association with execution of the write request, update the information indicating extent of page utilization of the specific one of the subdivisions;

[1e] logic to send the host the information indicating extent of page utilization of the specific one of the subdivisions and an address corresponding to the specific one of the subdivisions;

[1f] logic to receive an erasure request from the host, via the host interface, the erasure request accompanied by the address corresponding to the specific one of the subdivisions, to responsively control erasure of each unerased erase block in the group respective to the specific one of the subdivisions, and to update the information indicating extent of page utilization of the specific one of the subdivisions; and

[1g] logic to, in connection with the erasure, detect a defect in one of the erase blocks in the group respective to the specific one of the subdivisions,

to responsively substitute one or more different erase blocks of the flash memory for the one of the erase blocks for which the defect was detected;

[1h] wherein each said logic comprises at least one of hardware circuitry or instructions stored on non-transitory, machine-readable media that when executed are to control the operation of hardware circuitry.

746. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

747. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

748. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

749. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

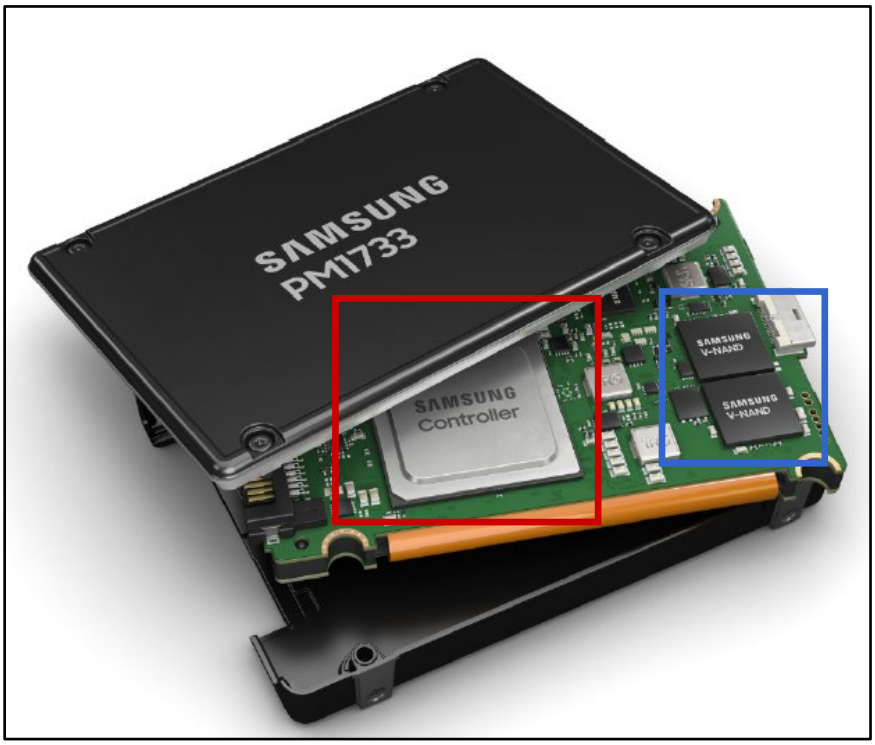
750. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

751. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

752. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a memory controller to control flash memory, the flash memory having subdivisions, each of the subdivisions corresponding to a respective group of erase blocks, each of the erase blocks such that it must be erased before storage locations in that erase block can be

reprogrammed. '183 Patent, claim 1, preamble.

753. For example, Samsung's PM1733 contains a memory controller (annotated in red) which controls flash memory chips (annotated in blue). See:



Samsung PM1733 Product Brief[33]

754. The flash memory in Samsung's Accused Instrumentalities contains erase blocks, which must be erased before storage locations in that erase block can be reprogrammed. See:

[33] Ex. Q, Samsung PM1733 Product Brief (annotated).

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Samsung Over-Provisioning White Paper[34]

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[35]

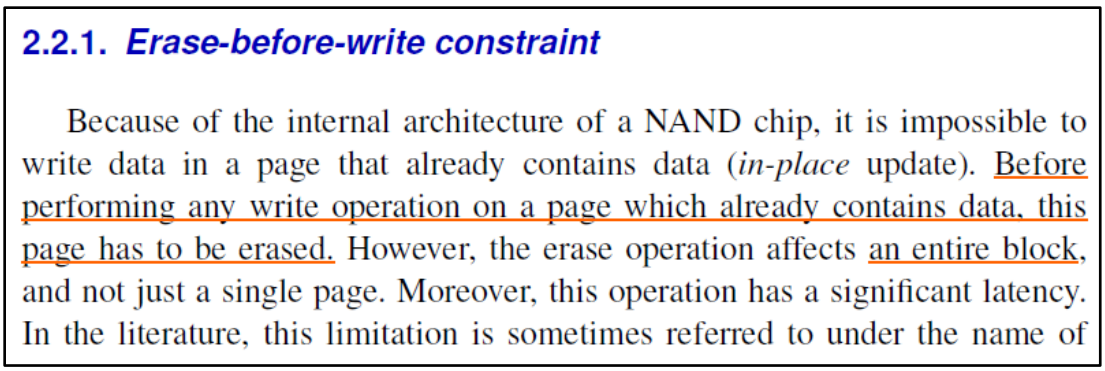
– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[36]

[34] Ex. V, Samsung, *Over-Provisioning Benefits for Samsung Data Center SSDs*, available at: https://download.semiconductor.samsung.com/resources/white-paper/S190311-SAMSUNG-Memory-Over-Provisioning-White-paper_10129514063996.pdf, at 2 (the "Samsung Over-Provisioning White Paper") (highlighted).

[35] Ex. I, Samsung SSD White Paper, at 17.

[36] Ex. W, Jalil Boukhobza et al., FLASH MEMORY INTEGRATION (ISTE Press Ltd 2017) 22 (highlighted).

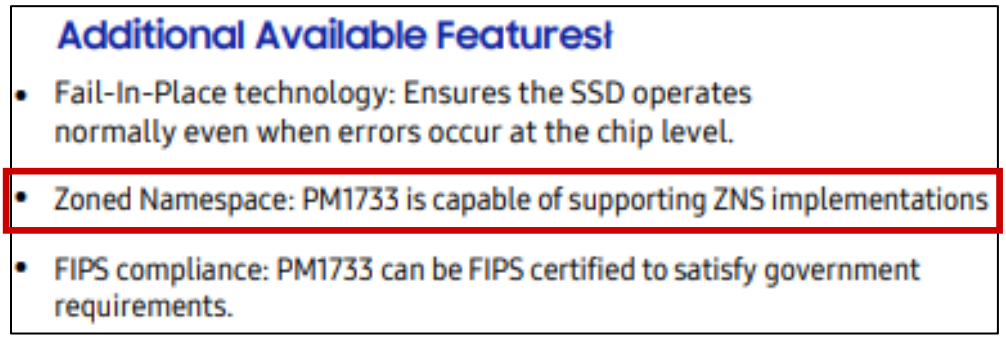


2.2.1. *Erase-before-write constraint*

Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[37]

755. The flash memory in Samsung's Accused Instrumentalities contains zones (subdivisions), each zone corresponding to a group of erase blocks. See:



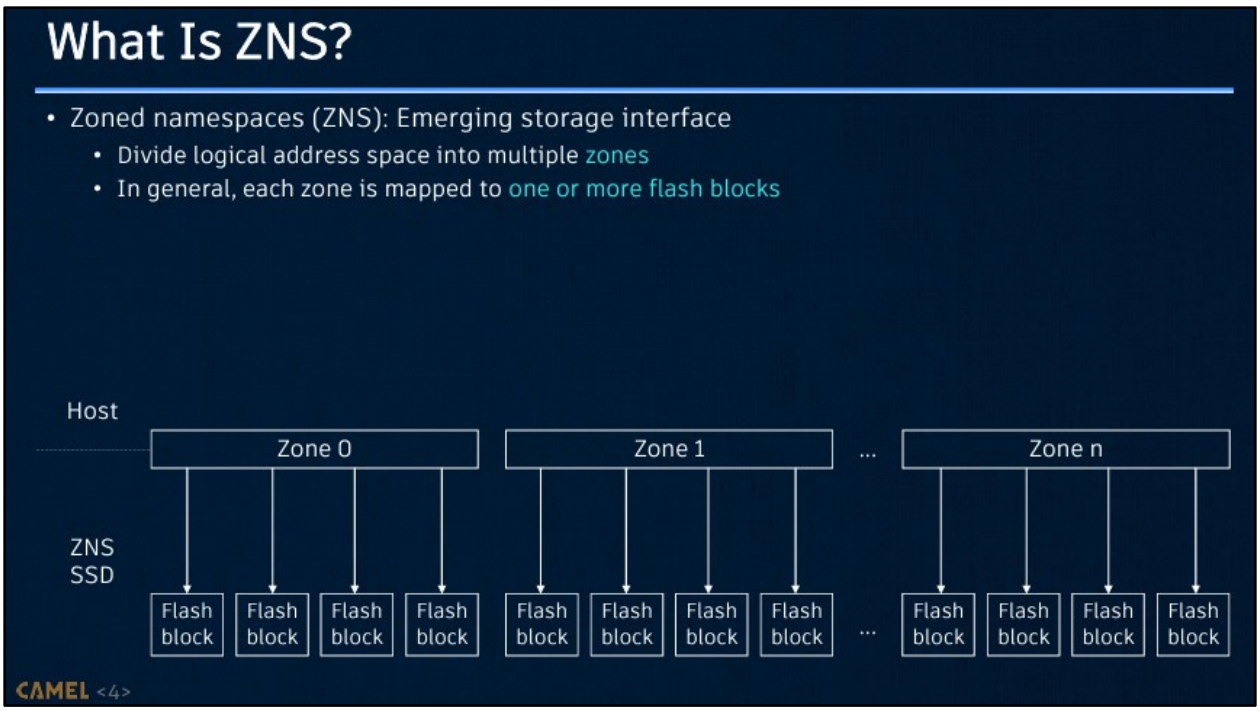
Additional Available Features!

- Fail-In-Place technology: Ensures the SSD operates normally even when errors occur at the chip level.
- Zoned Namespace: PM1733 is capable of supporting ZNS implementations
- FIPS compliance: PM1733 can be FIPS certified to satisfy government requirements.

Samsung PM1733 Product Brief[38]

[37] *Id*. at 24.

[38] Ex. Q, Samsung PM1733 Product Brief (annotated).



Presentation on ZNS[39]

> This new attribute was introduced to allow for the zone size to remain a power of two number of logical blocks (facilitating logical block to zone number conversions) while allowing optimized mapping of a zone storage capacity to the underlying media characteristics. For instance, in the case a flash based device, a zone capacity can be aligned to the size of flash erase blocks without requiring that the device implements a power-of-two sized erased block.

SSDs with NVMe Zoned Namespace (ZNS) Support[40]

---

[39] Ex. X, Hanyeoreum Bae et al., *What You Can't Forget: Exploiting Parallelism for Zoned Namespaces*, available at: https://www.hotstorage.org/2022/slides/hotstorage22-paper24-presentation_slides.pdf, at 4.

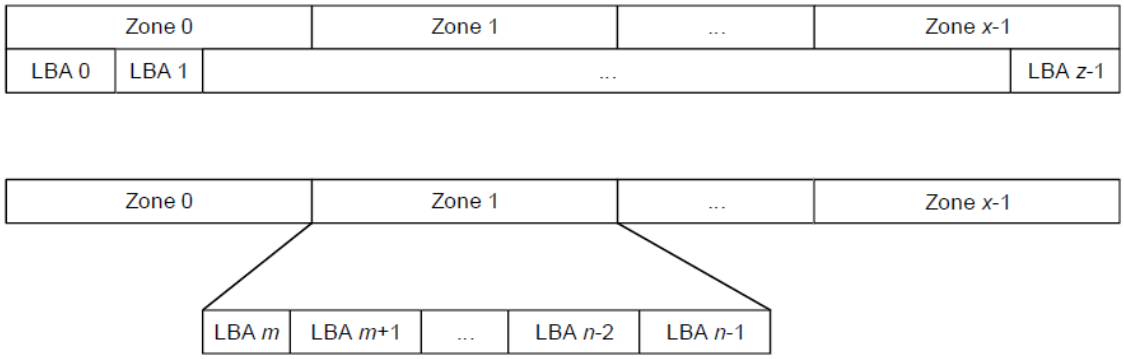
[40] Ex. Y, ZONED STORAGE, *SSDs with NVMe Zoned Namespace (ZNS) Support*, available at: https://zonedstorage.io/docs/introduction/zns.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[41]

756. Samsung's Accused Instrumentalities contain a host interface. '183 Patent, claim 1, element [1a].

757. For example, Samsung's PM1733 contains a PCIe host interface. See:

[41] Ex. Z, NVM EXPRESS, *Zoned Namespace Command Set Specification*, Revision 1.1d, (the "ZNS Specification") at 8.

Samsung PM1733 specifications	
Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[42]

758. Samsung's Accused Instrumentalities contain a memory interface. '183 Patent, claim 1, element [1b].

759. For example, Samsung's PM1733 contains a Toggle memory interface for communication with the NAND flash memory. See:

Samsung PM1733 specifications	
Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[43]

[42] Ex. Q, Samsung PM1733 Product Brief (annotated).

[43] *Id.*

Toggle NAND vs. ONFI NAND

In the NAND flash industry, there is some variation in implementation depending on manufacturer. Samsung and Toshiba make Toggle NAND, whereas everyone else makes what is called ONFI NAND.
Each generation of NAND, whether Toggle or ONFI, differs in the speeds it is capable of delivering: ONFI 1.0 and pre-Toggle NAND from Samsung and Toshiba (both asynchronous) are capable of speeds up to 50MB/s, ONFI 2.x (synchronous) and Toggle 1.0 NAND (asynchronous) are capable of delivering up to 200MB/s and 133MB/s respectively, and ONFI 3.0 and Toggle 2.0 (both asynchronous) are capable of delivering up to 400MB/s. Today's newest NAND chips are either Toggle 2.0 or ONFI 3.x.

Samsung SSD White Paper[44]

760. Samsung's Accused Instrumentalities contain logic to track information indicating extent of page utilization of each of the subdivisions. '183 Patent, claim 1, element [1c].

761. Samsung's Accused Instrumentalities are able to maintain a write pointer for each zone. The write pointer indicates extent of page utilization of the zone. See:

Write Pointer	The Write Pointer attribute defines the lowest numbered writeable logical block address in that zone. The validity of the write pointer is zone state specific and is defined per zone type (refer to section 2.1.1.2).

ZNS Specification, Section 2.1.1.1[45]

2.1.1.2.1.1 Writing in Sequential Write Required Zones

The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:

- Write command;
- Write Zeroes command;
- Write Uncorrectable command;
- Copy command; and
- Zone Append command.

A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

[44] Ex. I, Samsung SSD White Paper at 15.

[45] Ex. Z, ZNS Specification, at 9.

ZNS Specification, Section 2.1.1.2.1.1[46]

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

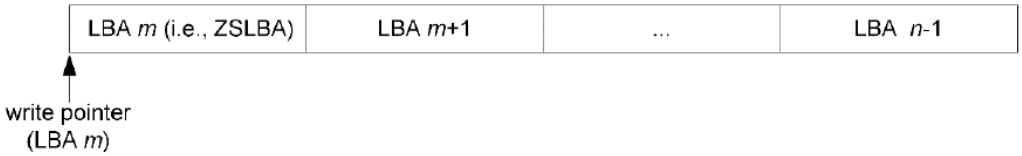
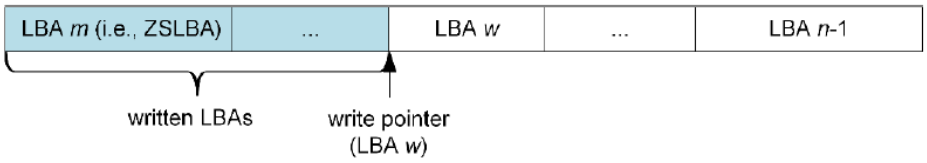


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[47]

762. Samsung's Accused Instrumentalities contain logic to receive a write request from a host via the host interface, the write request accompanied by an address that designates a specific one of the subdivisions, wherein the logic to track is to, in association with execution of the write request, update the information indicating extent of page utilization of the specific one of the subdivisions. '183 Patent, claim 1, element [1d].

763. Samsung's Accused Instrumentalities are able to receive and execute a Write command from a host.

---

[46] *Id*.

[47] *Id*. at 10.

764. A Write command is a write request.

765. A Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. See:

> **3.3.6 Write command**
>
> The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

NVM Command Set Specification, Section 3.3.6[48]

766. The Write command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:

[48] Ex. AA, NVM EXPRESS, *NVM Command Set Specification*, Revision 1.1 (the "NVM Command Set Specification"), at 53.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[49]

[49] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[50]

767. The Write command is accompanied by a logical block address ("LBA"). See:

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2

[50] Ex. BB, NVM EXPRESS, NVM Express Base Specification, Revision 2.0d (the "NVMe Base Specification"), at 107.

ZNS Specification, Section 3.2[51]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[52]

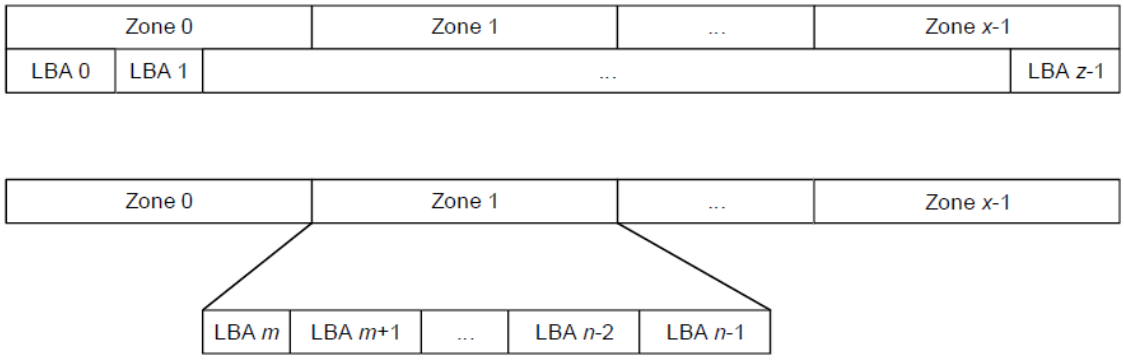
768. The LBA accompanying a Write command corresponds to a specific one of the zones. See:

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

[51] Ex. Z, ZNS Specification, at 18 (annotated).

[52] Ex. AA, NVM Command Set Specification, at 54.

ZNS Specification, Section 2.1.1[53]

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[54]

769. Samsung's Accused Instrumentalities are also able to receive and execute a Zone Append command from a host.

770. A Zone Append command is a write request.

771. The Zone Append command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:

[53] Ex. Z, ZNS Specification, at 8.

[54] *Id*. at 20.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[55]

[55] *Id.* at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[56]

772. The Zone Append command is accompanied by a Zone Start Logical Block Address ("ZSLBA") that designates a specific one of the zones. See:

3.4 Zoned Namespace Command Set I/O Commands

3.4.1 Zone Append command

The Zone Append command writes data and metadata, if applicable, to the I/O controller for the zone indicated by the ZSLBA field. The controller assigns the data and metadata, if applicable, to a set of logical blocks within the zone. The lowest LBA of the set of logical blocks written is returned in the completion queue entry (refer to section 3.4.1.2). The host may also specify protection information to include as part of the operation.

This command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used. All other command specific fields are reserved.

[56] Ex. BB, NVMe Base Specification, at 107.

Figure 25: Zone Append – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Zone Start Logical Block Address (ZSLBA):** This field indicates the 64-bit address of the lowest logical block of the zone in which the data and metadata, if applicable, associated with this command is to be stored. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

ZNS Specification, Section 3.4.1[57]

773. Samsung's Accused Instrumentalities are able to update the write pointer for a zone in association with execution of a write request directed to that zone. See:

2.1.1.2.1.1 Writing in Sequential Write Required Zones

The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:

- Write command;
- Write Zeroes command;
- Write Uncorrectable command;
- Copy command; and
- Zone Append command.

A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

ZNS Specification, Section 2.1.1.2.1.1[58]

The write pointer for a zone in the ZSE:Empty state, the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state shall be increased by the number of logical blocks written on successful completion of a write operation.

ZNS Specification, Section 2.1.1.2.1.1[59]

[57] Ex. Z, ZNS Specification, at 21–22.

[58] *Id*. at 9.

[59] *Id*. at 10.

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

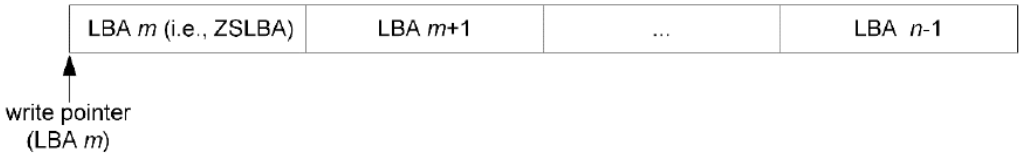
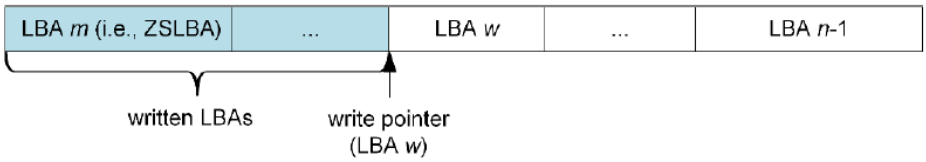


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[60]

774. Samsung's Accused Instrumentalities contain logic to send the host the information indicating extent of page utilization of the specific one of the subdivisions and an address corresponding to the specific one of the subdivisions. '183 Patent, claim 1, element [1e].

775. Samsung's Accused Instrumentalities are able to send the host the write pointer and the starting LBA of a zone, for example, in response to a Zone Management Receive command specifying a Report Zones action. See:

[60] *Id*. at 10.

3.4.2.1 Zone Receive Actions

The Zone Management Receive command Zone Receive Action field specifies what action to perform.

3.4.2.1.1 Report Zones

The Report Zones action returns the Report Zones data structure (refer to Figure 35). The Zone Descriptors of the Report Zones data structure shall:

a) report only Zone Descriptors of zones for which the ZSLBA value is greater than or equal to the ZSLBA value of the zone specified by the SLBA value in the command;
b) match the criteria in the Zone Receive Action Specific field; and
c) be sorted in ascending order by the ZSLBA value of each zone.

ZNS Specification, Section 3.4.2.1[61]

3.4.2.2.1 Report Zones Data Structure

Figure 35 defines the Report Zones Data Structure.

Figure 35: Report Zones Data Structure

Bytes	Description
07:00	**Number of Zones:** If the Partial Report bit (refer to Figure 34) is cleared to '0', then this field indicates the number of zones that match the criteria defined in section 3.4.2.1.1. If the Partial Report bit is set to '1', then this field indicates the number of zones for which complete Zone Descriptors were transferred to the data buffer. Refer to section 3.4.2.1.1 for the content of the data buffer.
63:08	Reserved
127:64	**Zone Descriptor 0:** Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).
191:128	**Zone Descriptor 1:** Contains the Zone Descriptor for the second zone reported, if any.
...	...
((n+1)*64)+63:(n+1)*64	**Zone Descriptor n:** Contains the Zone Descriptor for the last zone reported, if any.

ZNS Specification, Section 3.4.2.2.1[62]

[61] *Id*. at 26.

[62] *Id*.

3.4.2.2.3 Zone Descriptor Data Structure

Figure 37 defines the Zone Descriptor data structure.

Figure 37: Zone Descriptor Data Structure

Bytes	Description	
	Bits	**Description**
00	7:4	Reserved
	3:0	**Zone Type (ZT):** This field indicates the type of the zone. **Value** / **Definition** 2h / Sequential Write Required All other values / Reserved
23:16		**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.
31:24		**Write Pointer (WP):** This field is the logical block address where the next write operation for this zone should be issued. Refer to section 2.1.1.2.1 for the behavior of the write pointer.

ZNS Specification, Section 3.4.2.2.3[63]

776. Samsung's Accused Instrumentalities contain logic to receive an erasure request from the host, via the host interface, the erasure request accompanied by the address corresponding to the specific one of the subdivisions, to responsively control erasure of each unerased erase block in the group respective to the specific one of the subdivisions, and to update the information indicating extent of page utilization of the specific one of the subdivisions. '183 Patent, claim 1, element [1f].

777. Samsung's Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action from the host. See:

[63] *Id*. at 27–28.

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

ZNS Specification, Section 3.4.3[64]

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

[64] *Id*. at 29.

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[65]

778. A Zone Management Send command specifying a Reset Zone action is an erasure request.

779. The Zone Management Send command is accompanied by the starting LBA of a zone. See:

3.4.3 Zone Management Send command

The Zone Management Send command requests an action on one or more zones. The command uses the Data Pointer, Command Dword 10, Command Dword 11, and Command Dword 13 fields. All other command specific fields are reserved.

Figure 38: Zone Management Send – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field specifies the lowest LBA of the zone on which the Zone Send Action is performed. Command Dword 10 contains bits 31:00 of the SLBA; Command Dword 11 contains bits 63:32 of the SLBA.

[65] *Id*. at 30–31.

ZNS Specification, Section 3.4.3[66]

780. Samsung's Accused Instrumentalities are able to responsively control erasure of each unerased erase block in the zone. See:

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[67]

781. Samsung's Accused Instrumentalities are able to update the write pointer of the zone in response to a Zone Management Send command specifying a Reset Zone action. See:

[66] *Id*. at 29.

[67] *Id*. at 30–31.

> The Zone Management Send command with a Zone Send Action of Reset Zone sets the write pointer to the ZSLBA for that zone.

ZNS Specification, Section 2.1.1.2.1.1[68]

> If the command completes successfully, then for each affected zone:
>
> a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone;

ZNS Specification, Section 3.4.3.1.4[69]

782. Samsung's Accused Instrumentalities contain logic to, in connection with the erasure, detect a defect in one of the erase blocks in the group respective to the specific one of the subdivisions, to responsively substitute one or more different erase blocks of the flash memory for the one of the erase blocks for which the defect was detected. '183 Patent, claim 1, element [1g].

783. Samsung's Accused Instrumentalities are able to detect bad erase blocks in connection with erasure, map out the bad erase blocks, and substitute different erase blocks for the bad ones. See:

> **Bad Block Management & Error Correcting Code (ECC)**
>
> In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

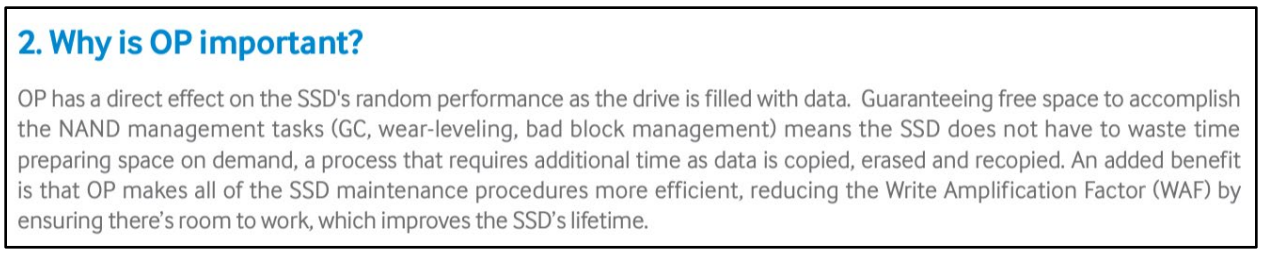
Samsung SSD White Paper[70]

---

[68] *Id*. at 10.

[69] *Id*. at 30–31.

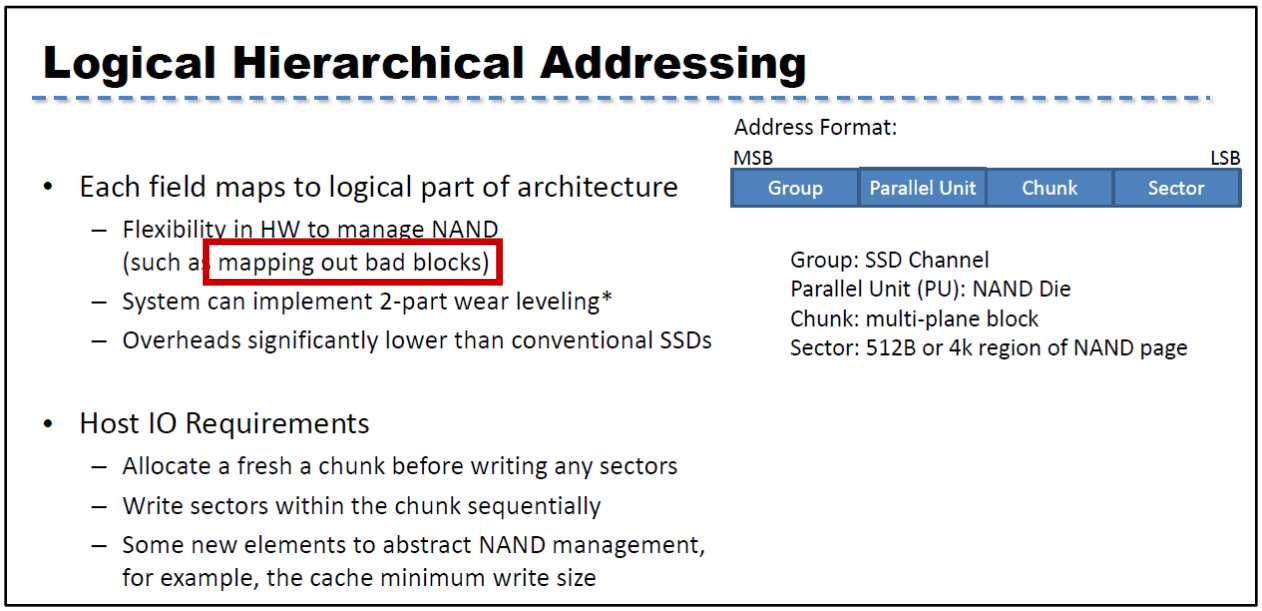
[70] Ex. I, Samsung SSD White Paper, at 19.



2. Why is OP important?

OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

Samsung Application Note[71]



Microsoft Denali Presentation[72]

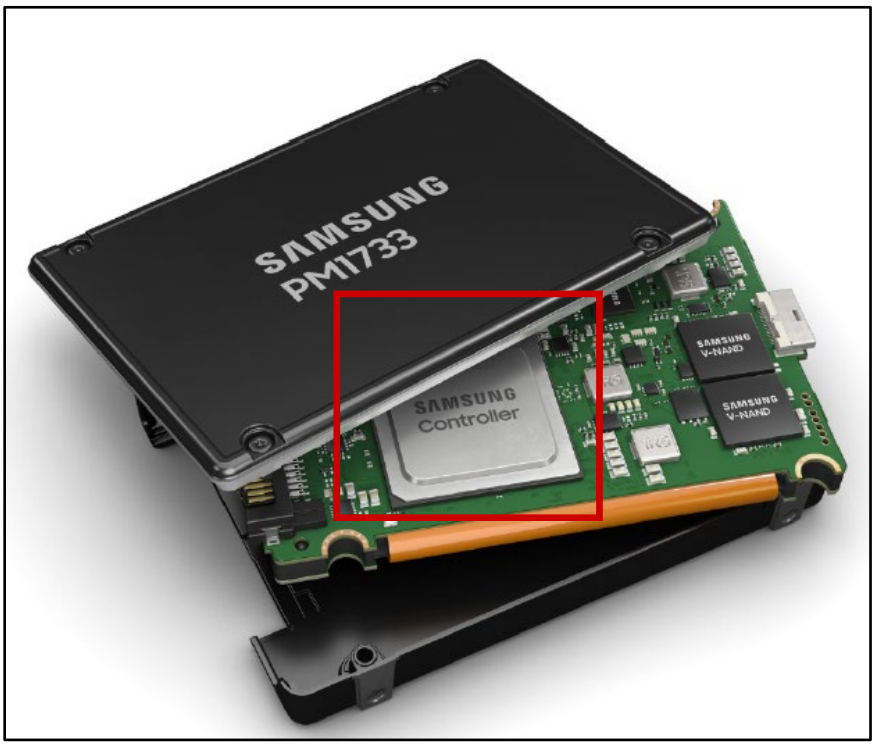
784. In Samsung's Accused Instrumentalities, each said logic comprises at least one of hardware circuitry or instructions stored on non-transitory, machine-readable media that when executed are to control the operation of hardware circuitry. '183 Patent, claim 1, element [1h]. See,

---

[71] Ex. CC, SAMSUNG, *Over-provisioning Maximize the Lifetime and Performance of Your SSD with Small Effect to Earn More*, available at: https://download.semiconductor.samsung.com/resources/others/Samsung_SSD_845DC_04_Over-provisioning.pdf, at 2 (the "Samsung Application Note.")

[72] Ex. DD, MICROSOFT, Denali the Next Generation High Density Storage Interface, March 2018 OCP Summit, at 15 (the "Microsoft Denali Presentation") (annotated).

for example:



Samsung PM1733 Product Brief[73]

785. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '183 Patent.

786. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

787. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

788. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

---

[73] Ex. Q, Samsung PM1733 Product Brief (annotated).

789. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

790. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

791. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

792. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

793. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

794. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

795. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

796. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

797. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

798. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

799. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

800. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

801. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

802. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

803. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

804. SEC induces infringement by Samsung's customers by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

805. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

806. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

807. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or

facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

808. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

809. SEC contributes to infringement by other Samsung subsidiaries by offering to sell,

selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

810. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

811. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

812. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

813. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building,

manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

814. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

815. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or

hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

816. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

817. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

818. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

819. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

820. SEC's infringement of the Asserted Patent is ongoing.

821. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

822. SEC's ongoing infringement of the Asserted Patent is willful.

823. SEC's ongoing infringement of the Asserted Patent is egregious.

824. SEA's infringement of the Asserted Patent is ongoing.

825. SEA will not stop its infringement of the Asserted Patent absent a Court order

restraining it from future infringement.

826. SEA's ongoing infringement of the Asserted Patent is willful.

827. SEA's ongoing infringement of the Asserted Patent is egregious.

828. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

829. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

830. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

831. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

832. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

833. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT TWO: INFRINGEMENT OF THE '772 PATENT

834. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

835. U.S. Patent No. 11,709,772 ("the '772 Patent"), entitled "Storage System with Multiplane Segments and Cooperative Flash Management," was legally and duly issued on July 25, 2023, naming Andrey V. Kuzmin and James G. Wayda as the inventor. *See* Exhibit B.

836. Radian owns all rights, title, and interest in the '772 Patent, and holds all substantial

rights pertinent to this suit, including the right to sue and recover for all past, current, and future infringement.

837. The '772 Patent is valid, enforceable, and directed to patentable subject matter.

838. Radian has complied with 35 U.S.C. § 287 with respect to the '772 Patent.

839. The '772 Patent described challenges faced by conventional memory devices, including the difficulties to realize the multi-plane benefits. See, for example:

> Unfortunately, it is difficult to realize the benefits of multi-plane or multi-die architecture at an application or operating system level. That is to say, logical addresses typically arrive at the memory controller from the host as a stream of random accesses, and are sequentially assigned to first available physical space by the memory controller; as memory, particularly flash memory, is erased and recycled via read, write, wear leveling, garbage collection and other processes, sequential logical addresses become scattered throughout physical memory space. For multi-plane memory therefore, there is no practical mechanism for the host or memory controller to group related data in a manner geared for multi-plane access (i.e., there is no guarantee available physical addresses used for related data will be consistent with device multi-plane addressing restrictions).

> A need therefore exists for improvements that provide additional flexibility to systems and application designers. Ideally, such improvements would provide flexibility in storing and accessing multiple pages of data, for example, across multiple dies or planes.[74]

840. The '772 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> [W]ith insight into geometry, a host can organize address space in a manner geared for improved memory access, e.g.,

[74] Ex. B, U.S. Patent No. 11,709,772, at 2:5–41.

> according to a policy aimed at IO optimization for managing multiple or linked pages.[75]
>
> In applying these techniques, where allocation requests allow, the host assigns physical space for writes so as to facilitate both immediate write accesses and later multiple page read access, all with the benefit of the improved throughput. For example, later accesses can be performed in a manner unencumbered by memory controller address translation requirements; for multi-plane memory and shingled drives, data placement can be pre-planned in a manner consistent with any inter-plane addressing restrictions for multi-array accesses.[76]
>
> Optional features mentioned above are supported by structure where the memory controller also stores information specific to each of plural subdivisions of memory (e.g., for each physical memory unit managed by the host, in on-board registers reserved for this purpose). The memory controller makes data based on this stored information accessible to the host.[77]

841. Some of these described solutions are reflected in the claimed inventions of the ’772 Patent.

842. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the ’772 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 2 of the ’772 Patent.

843. Claim 2 of the ’772 Patent recites:

> [preamble] A solid state storage drive (SSD) comprising:

---

[75] *Id*. at 4:61–64.

[76] *Id*. at 5:1–10.

[77] *Id*. at 12:3–9.

[2a] a host interface configured to communicate with a host using commands from a Non-Volatile Memory Express (NVMe) standard;

[2b] non-volatile memory comprising NAND flash memory, the NAND flash memory comprising a plurality of erase units and each erase unit comprising a plurality of pages;

[2c] logic coupled to the host interface and to the flash memory, wherein the logic comprises at least one of hardware logic and software logic, and communicates with the host via the host interface using NVMe commands;

[2d] wherein the logic operates to:

[2e] divide the flash memory into zones; wherein the zones are non-overlapping subdivisions of the flash memory, each zone comprises two or more erase units, the erase units in at least one zone are on two or more different planes, and each zone is associated with a zone start logical block address and a range of logical block addresses;

[2f] in response to a write command from the host to write data within the range of logical block addresses of a specified zone, perform a multi-plane write of said data to erase units on different planes in the specified zone using a common page address for the respective erase units on the different planes;

[2g] generate metadata based on accessing and maintenance of erase units in the specified zone; and

[2h] based on the metadata, make a notification available to the host that recommends performing maintenance on the specified zone,

wherein the recommendation is associated with the zone start logical block address of the specified zone.

844. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

845. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

846. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

847. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

848. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

849. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

850. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a solid state storage drive (SSD). '772 Patent, claim 2, preamble.

851. For example, Samsung's PM1733 is an SSD. See:

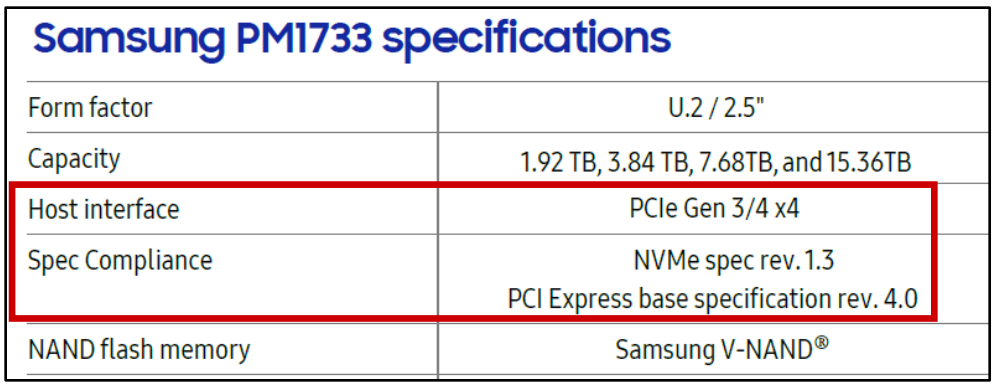


Samsung PM1733 Product Brief[78]

852. Samsung's Accused Instrumentalities contain a host interface configured to communicate with a host using commands from a Non-Volatile Memory Express (NVMe) standard. '772 Patent, claim 2, element [2a].

853. For example, Samsung's PM1733 contains a host interface configured to communicate with a host using NVMe commands. See:

[78] Ex. Q, Samsung PM1733 Product Brief.



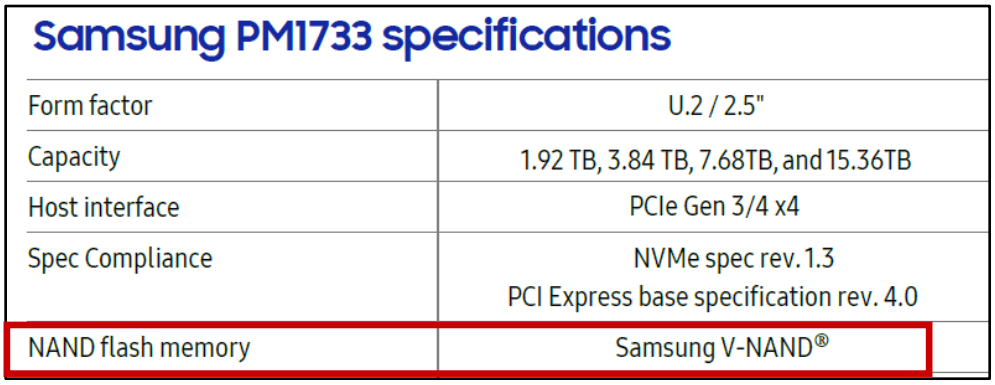
Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[79]

854. Samsung's Accused Instrumentalities contain non-volatile memory comprising NAND flash memory, the NAND flash memory comprising a plurality of erase units and each erase unit comprising a plurality of pages. '772 Patent, claim 2, element [2b].

855. For example, Samsung's PM1733 contains non-volatile memory comprising NAND flash memory. See:



Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

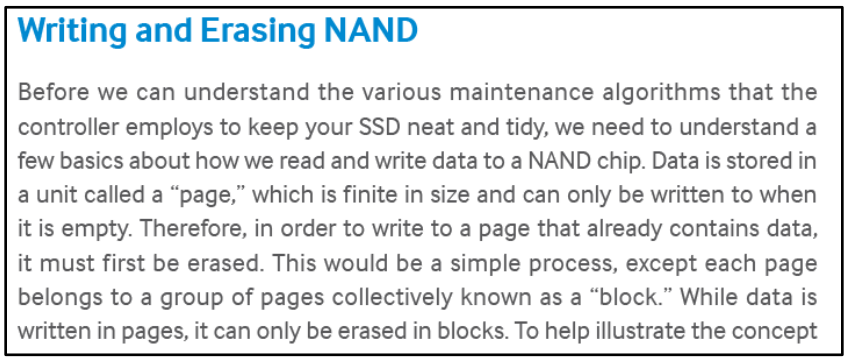
Samsung PM1733 Product Brief[80]

856. The NAND flash memory in Samsung's Accused Instrumentalities comprises a

[79] *Id*. (annotated).

[80] *Id*. (annotated).

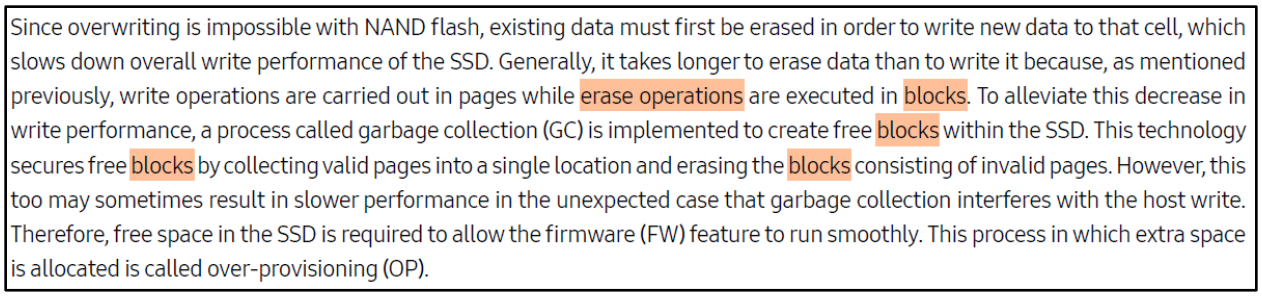
plurality of blocks (erase units). See:



Writing and Erasing NAND

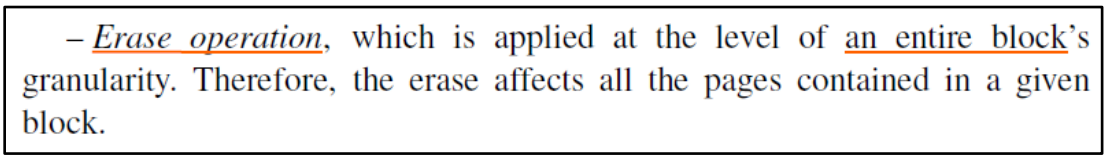
Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[81]



Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Samsung Over-Provisioning White Paper[82]



– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[83]

[81] Ex. I, Samsung SSD White Paper, at 17.

[82] Ex. V, Samsung Over-Provisioning White Paper, at 2 (highlighted).

[83] Ex. W, FLASH MEMORY INTEGRATION, at 22 (highlighted).

2.2.1. *Erase-before-write constraint*

Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[84]

857. Each block comprises a plurality of pages. See:

Writing and Erasing NAND

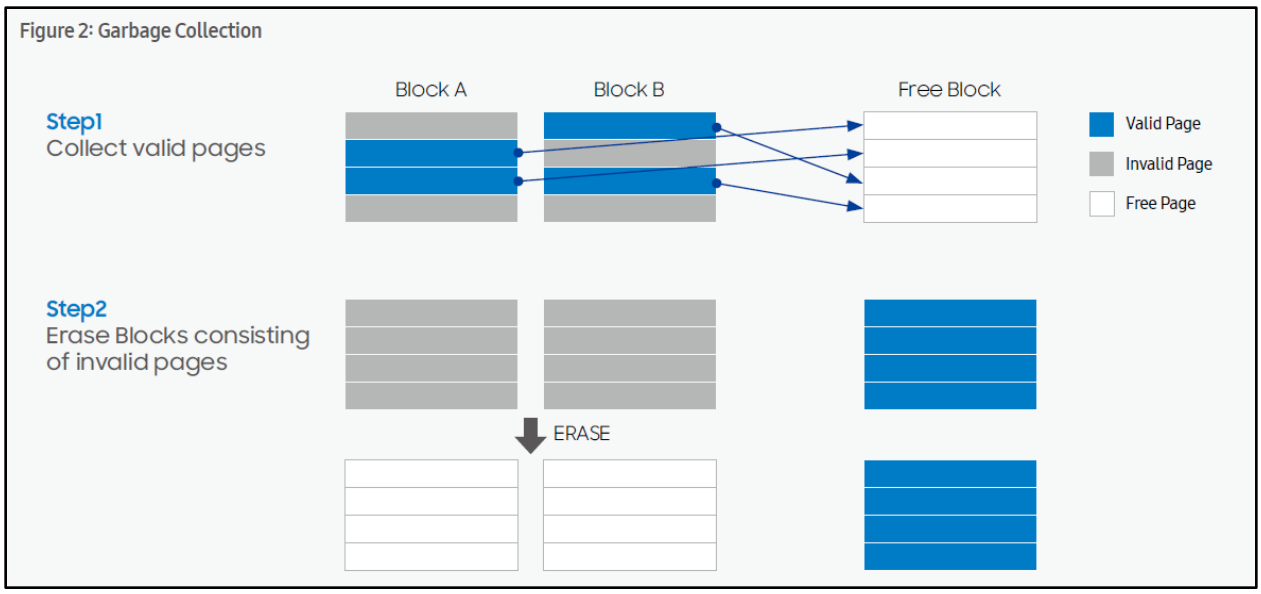
Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[85]

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

[84] *Id*. at 24.

[85] Ex. I, Samsung SSD White Paper, at 17.



Samsung Over-Provisioning White Paper[86]

[86] Ex. V, Samsung Over-Provisioning White Paper, at 2.

2.1.3. *Simplified hierarchical architecture of a NAND flash memory chip*

Figure 2.2 illustrates a simplified architecture of a NAND flash memory chip. This is a high-level vision, relatively abstract with respect to the micro-architectural level. A flash memory chip has a hierarchical structure. It is composed of a certain number of *planes*. Planes contain *blocks*, and blocks contain *pages*. The chip is equipped with an input/output (I/O) bus which makes the reception of commands and addresses possible, as well as the transfer of data to or from the host computer system. This bus is usually multiplexed, which means that commands, addresses and data share the same pins of the flash chip. Depending on the chip version, buses with 8 or 16 bits can be found.

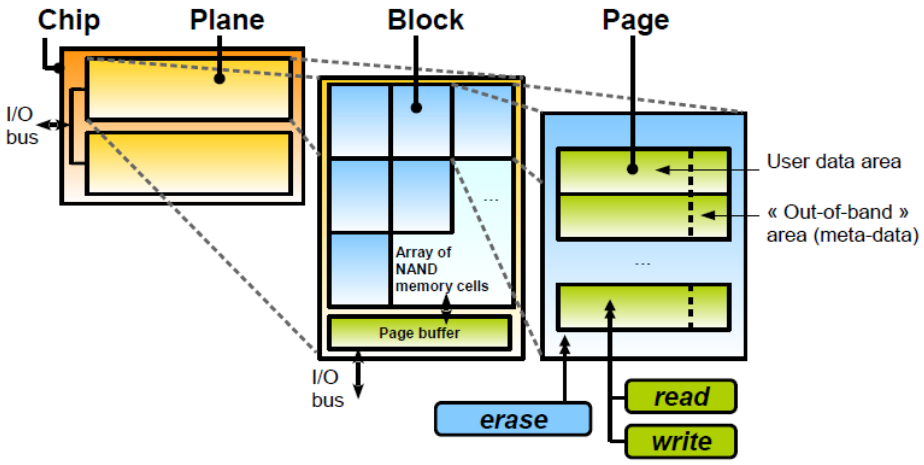


Figure 2.2. *Simplified architecture of a NAND flash memory chip. For a color version of this figure, see www.iste.co.uk/boukhobza/flash.zip*

Flash Memory Integration[87]

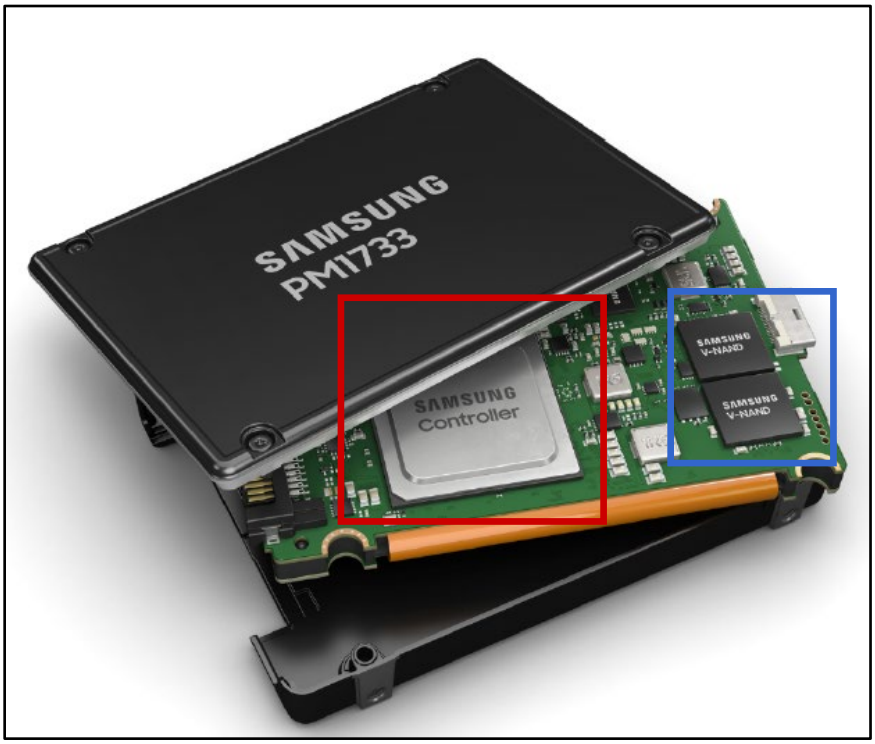
– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

[87] Ex. W, FLASH MEMORY INTEGRATION, at 18.

Flash Memory Integration[88]

858. Samsung's Accused Instrumentalities contain logic coupled to the host interface and to the flash memory, wherein the logic comprises at least one of hardware logic and software logic, and communicates with the host via the host interface using NVMe commands. '772 Patent, claim 2, element [2c].

859. Samsung's Accused Instrumentalities contain a host interface and flash memory. *See* '772 Patent, claim 2, elements [2a], [2b]. Samsung's Accused Instrumentalities contain logic (controller annotated in red) coupled to the host interface and the flash memory (NAND annotated in blue) comprising at least one of hardware logic and software logic. See, for example:



[88] *Id*. at 22.

Samsung PM1733 Product Brief[89]

860. Samsung's Accused Instrumentalities contain logic which communicates with the host via the host interface using NVMe commands. See:

Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[90]

861. Samsung's Accused Instrumentalities satisfy the claim element "wherein the logic operates to." '772 Patent, claim 2, element [2d]. Samsung's Accused Instrumentalities contain logic that operates to perform elements [2e] through [2h]. See below.

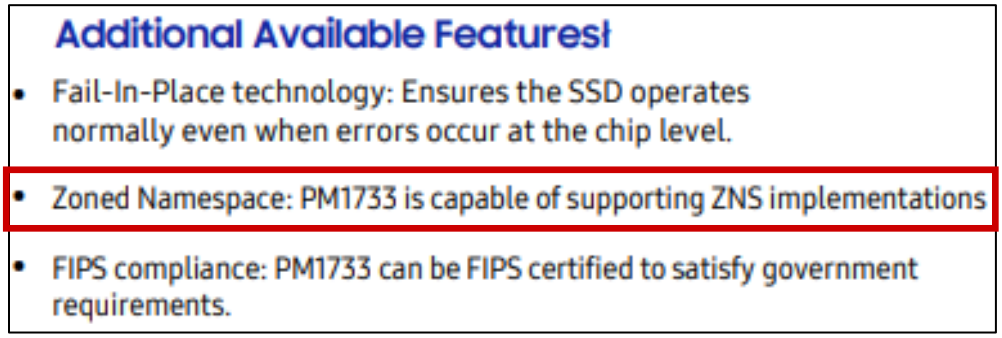
862. Samsung's Accused Instrumentalities contain logic that operates to divide the flash memory into zones; wherein the zones are non-overlapping subdivisions of the flash memory, each zone comprises two or more erase units, the erase units in at least one zone are on two or more different planes, and each zone is associated with a zone start logical block address and a range of logical block addresses. '772 Patent, claim 2, element [2e].

863. Samsung's Accused Instrumentalities are able to divide the flash memory into

[89] Ex. Q, Samsung PM1733 Product Brief (annotated).

[90] *Id*. (annotated).

zones. The zones are non-overlapping subdivisions of the flash memory. See:



Samsung PM1733 Product Brief[91]

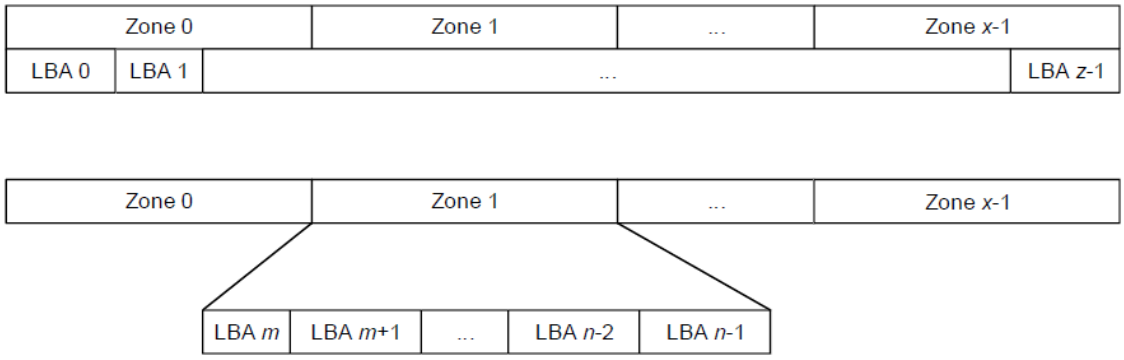
[91] *Id.* (annotated).

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



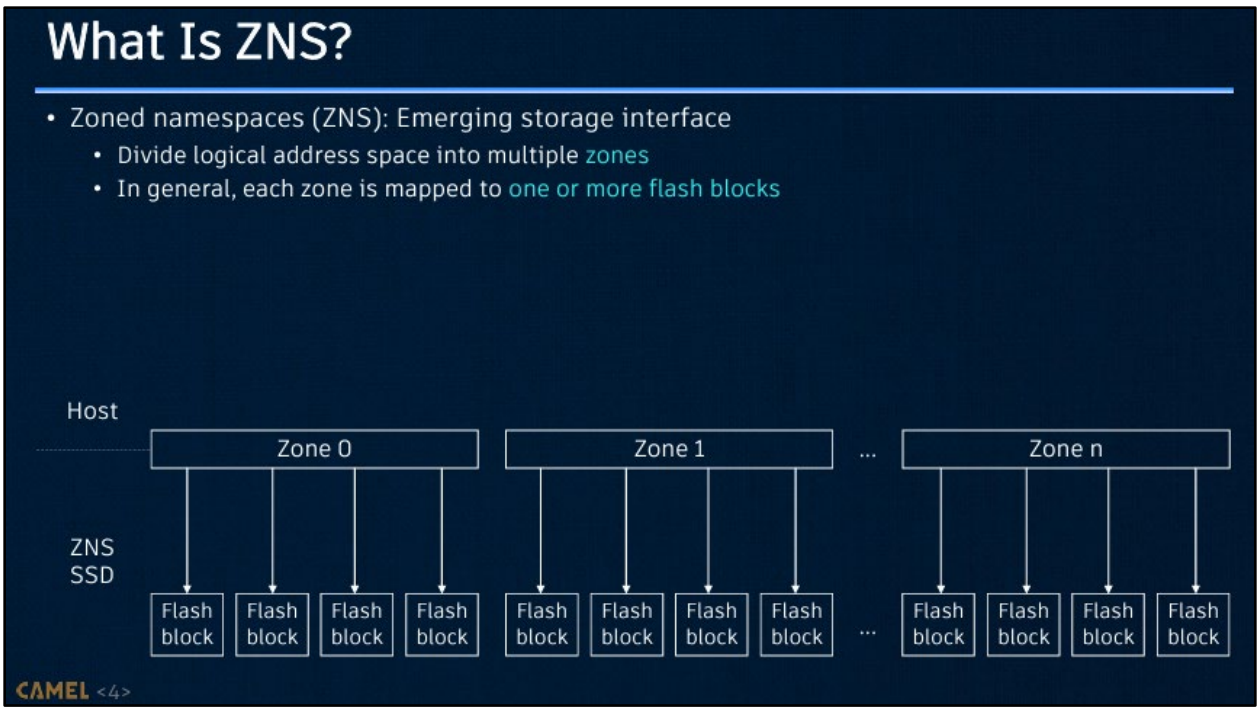
The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[92]

864. Each zone in Samsung's Accused Instrumentalities comprises two or more erase units. See:

[92] Ex. Z, ZNS Specification, at 8.



Presentation on ZNS[93]

> To maximize the flash operation parallelism, the flash blocks from a set of flash chips accessible in parallel will compose the FBG of a zone, and the chunks of a zone need to be interleavingly placed on the parallel flash chips. The number of parallel flash chips for the chunk interleaving is referred to as the *zone interleaving degree* D_{zone}, and the set of logically consecutive chunks across the parallel flash chips is referred to as a **stripe**. For a coarse-grained zone-to-FBG mapping, an FBG has flash blocks at the same block offset in the parallel flash chips, and the chunks of a stripe are located at the same page offset in different flash blocks.

[93] Ex. X, *What You Can't Forget: Exploiting Parallelism for Zoned Namespaces*, at 4.

Samsung Research Paper[94]

865. The erase units in at least one zone of Samsung's Accused Instrumentalities are on two or more different planes. See:

Memory density and die floor plan

Figure 1 shows 16 48L V-NAND dice with two F-Chips in an MCP (multichip package). Die efficiency is higher in the 48L. The 32L V-NAND die area is 84.3 mm^2 while the 48L V-NAND die measures 99.8 mm^2, which is a 17.3% increase due to greater die length (Figure 2). Memory density per unit die area is increased to 2.57 Gb/mm^2. The memory density of leading-edge 2D planar NAND devices such as Toshiba 15 nm TLC NAND is 1.28 Gb/mm^2, (see more on our analysis on the Toshiba 15 nm NAND). The key feature differences in die floor plans are 1) planar (NAND memory array) area, 2) bitline switch and page buffer area, 3) LOGIC and peripheral area, and 4) the addition of F-chip. **Each die has two planes.** The NAND memory array area increased from 48.9 mm^2 to 68.7 mm^2 which is 40.3% larger. The area of the bitline switch circuit is the same as that of the 32L, while the page buffer area is 20% smaller. The logic and peripheral circuit area is 34.8% reduced. In other words, Samsung dramatically shrank the area of the page buffer and peripheral region so that they can further increase memory density and die efficiency. Die thickness is also reduced from 132 μm to 36 μm for the sixteen die stacked in the MCP.

EETimes V-NAND analysis[95]

[94] Ex. EE, Kyuhwa Han et al., *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, 15th USENIX Symposium on Operating Systems Design and Implementation, at 149.

[95] Ex. FF, Jeongdong Choe, EETimes, *Samsung's 3D V-NAND 32L vs 48L–Just Vertical Expansion?*, available at: https://www.eetimes.com/samsungs-3d-v-nand-32l-vs-48l-just-vertical-expansion/ (highlighted).

Samsung TLC V- NAND Flash Memory		
	9th Gen V-NAND	8th Gen V-NAND
Layers	290?	236
Decks	2 (x145)	2 (x118)
Die Capacity	1 Tbit	1 Tbit
Die Size (mm^2)	?mm2	?mm2
Density ($Gbit/mm^2$)	?	?
I/O Speed	3.2 GT/s (Toggle 5.1)	2.4 GT/s (Toggle 5.0)
Planes	6?	4
CuA / PuC	Yes	Yes

AnandTech V-NAND analysis [96]

To maximize the flash operation parallelism, the flash blocks from a set of flash chips accessible in parallel will compose the FBG of a zone, and the chunks of a zone need to be interleavingly placed on the parallel flash chips. The number of parallel flash chips for the chunk interleaving is referred to as the *zone interleaving degree* D_{zone}, and the set of logically consecutive chunks across the parallel flash chips is referred to as a **stripe**. For a coarse-grained zone-to-FBG mapping, an FBG has flash blocks at the same block offset in the parallel flash chips, and the chunks of a stripe are located at the same page offset in different flash blocks.

Samsung Research Paper[97]

866. Each zone is associated with a zone start LBA and a range of LBA. See:

[96] Ex. GG, Anton Shilov, AnandTech, *Samsung Starts Mass Production of 9th Generation V-NAND: 1Tb 3D TLC NAND*, available at: https://www.anandtech.com/show/21365/samsung-starts-mass-production-of-9th-generation-vnand-1tb-3d-tlc-nand (annotated).

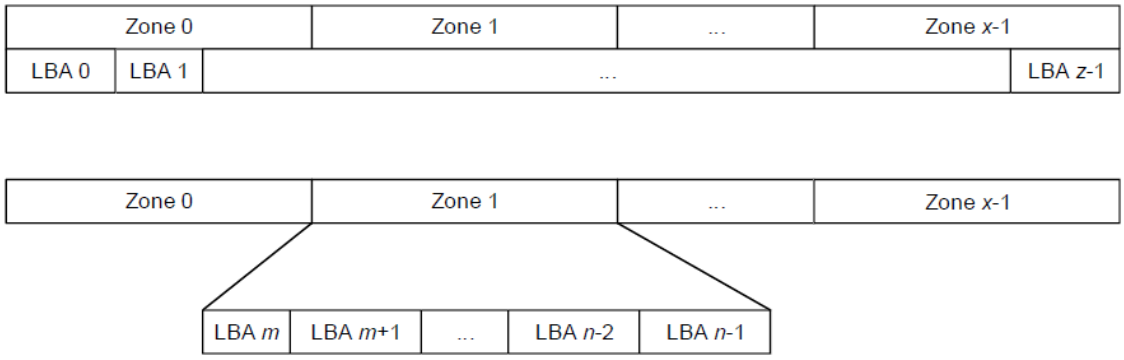
[97] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[98]

15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).
23:16	**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.

[98] Ex. Z, ZNS Specification, at 8.

ZNS Specification, Section 3.4.2.2.3[99]

63:00	**Zone Size (ZSZE):** This field contains the size of each zone in the zoned namespace. The value is reported as a number of logical blocks. The value of the field shall not be cleared to 0h.

ZNS Specification, Section 4.1.5.1[100]

867. Samsung's Accused Instrumentalities contain logic that operates to, in response to a write command from the host, write data within the range of logical block addresses of a specified zone, perform a multi-plane write of said data to erase units on different planes in the specified zone using a common page address for the respective erase units on the different planes. '772 Patent, claim 2, element [2f].

868. Samsung's Accused Instrumentalities are able to write data in response to a Write command from the host. The Write command specifies LBAs within the range of LBAs of a zone, otherwise the memory device will report an error. See:

[99] *Id*. at 28.

[100] *Id*. at 38.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[101]

869. Samsung's Accused Instrumentalities are also able to write data in response to a Zone Append command from the host. The Zone Append command specifies a zone, and the device writes data within the range of LBAs of the specified zone. See:

3.4.1 Zone Append command

The Zone Append command writes data and metadata, if applicable, to the I/O controller for the zone indicated by the ZSLBA field. The controller assigns the data and metadata, if applicable, to a set of logical blocks within the zone. The lowest LBA of the set of logical blocks written is returned in the completion queue entry (refer to section 3.4.1.2). The host may also specify protection information to include as part of the operation.

This command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used. All other command specific fields are reserved.

[101] *Id*. at 20 (annotated).

Figure 25: Zone Append – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Zone Start Logical Block Address (ZSLBA):** This field indicates the 64-bit address of the lowest logical block of the zone in which the data and metadata, if applicable, associated with this command is to be stored. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

ZNS Specification, Section 3.4.1[102]

870. Samsung's Accused Instrumentalities perform a multi-plane write to erase units on different planes in the specified zone using a common page address for the respective erase units on the different planes. See:

> To maximize the flash operation parallelism, the flash blocks from a set of flash chips accessible in parallel will compose the FBG of a zone, and the chunks of a zone need to be interleavingly placed on the parallel flash chips. The number of parallel flash chips for the chunk interleaving is referred to as the *zone interleaving degree* D_{zone}, and the set of logically consecutive chunks across the parallel flash chips is referred to as a **stripe**. For a coarse-grained zone-to-FBG mapping, an FBG has flash blocks at the same block offset in the parallel flash chips, and the chunks of a stripe are located at the same page offset in different flash blocks.

Samsung Research Paper[103]

871. Samsung's Accused Instrumentalities contain logic that operates to generate metadata based on accessing and maintenance of erase units in the specified zone. '772 Patent, claim 2, element [2g].

---

[102] *Id*. at 21–22.

[103] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

872. Samsung's Accused Instrumentalities generate metadata based on accessing and maintenance of erase units for purposes of internal operations. See:



Samsung Over-Provisioning White Paper[104]

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-

[104] Ex. CC, Samsung Over-Provisioning White Paper, at 1.

194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[105]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

[105] *Id.*, at 6.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

Samsung SSD White Paper[106]

2. Why is OP important?

OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

[106] Ex. I, Samsung SSD White Paper, at 19.

Samsung Application Note[107]

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[108]

873. Samsung's Accused Instrumentalities contain logic that operates to, based on the metadata, make a notification available to the host that recommends performing maintenance on the specified zone, wherein the recommendation is associated with the zone start logical block address of the specified zone. '772 Patent, claim 2, element [2h].

874. ZNS Specification defines a "Reset Zone Recommended" attribute, which is a notification available to the host that recommends performing maintenance on a specified zone. See:

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

[107] Ex. CC, Samsung Application Note, at 2.

[108] Ex. Z, ZNS Specification, at 41.

to section 3.4.3.1.4). If a controller schedules such an internal operation on a zone, the controller may notify the host by:

a) setting the Reset Zone Recommended zone attribute of the specific zone to '1' (refer to Figure 37);
b) setting the Reset Zone Recommended Time Limit zone attribute information to indicate the number of seconds before the controller intends to perform an internal operation on the specified zone; and
c) generating a Zone Descriptor Changed event for the specific zone.

As a zone reset operation destroys data in a specific zone, it is optional for the host software to perform a zone reset operation on zones that have the Reset Zone Recommended zone attribute set to '1'. If the host does not perform a zone reset operation on the specific zone, then the internal operation, which may impact performance, may be performed.

If the controller has processed the internal operation or the internal operation is no longer scheduled, the controller may notify the host by:

a) clearing the Reset Zone Recommended zone attribute of the specific zone to '0'; and
b) generating a Zone Descriptor Changed event for the specific zone.

If a zone is in the ZSF:Read Only state or the ZSF:Offline state, then the Reset Zone Recommended attribute shall be cleared to '0'.

ZNS Specification, Section 5.4[109]

Figure 37: Zone Descriptor Data Structure

Bytes	Description
02	**Zone Attributes (ZA):** Indicates attributes for the Zone:

Bits	Description
7	**Zone Descriptor Extension Valid (ZDEV):** If this bit is set to '1', then Zone Descriptor Extension data is associated with the zone. If this bit is cleared to '0', then no Zone Descriptor Extension data is associated with the zone. Refer to section 5.1.
6:3	Reserved
2	**Reset Zone Recommended (RZR):** If this bit is set to '1', then the controller recommends that this zone be reset. Refer to section 5.4.
1	**Finish Zone Recommended (FZR):** If this bit is set to '1', then the controller recommends that this zone be finished. Refer to section 5.5.
0	**Zone Finished by Controller (ZFC):** If this bit is set to '1', then the controller finished this zone due to a Zone Active Excursion. Refer to section 5.6.

ZNS Specification, Section 3.4.2.2.3[110]

875. The Reset Zone Recommended attribute is associated with the zone start logical

[109] *Id*. at 41–42.

[110] *Id*. at 28.

block address (or ZSLBA) of the specified zone. See:

Figure 37: Zone Descriptor Data Structure

Bytes	Description
02	**Zone Attributes (ZA):** Indicates attributes for the Zone: **Bits** / **Description** 7 — **Zone Descriptor Extension Valid (ZDEV):** If this bit is set to '1', then Zone Descriptor Extension data is associated with the zone. If this bit is cleared to '0', then no Zone Descriptor Extension data is associated with the zone. Refer to section 5.1. 6:3 — Reserved 2 — **Reset Zone Recommended (RZR):** If this bit is set to '1', then the controller recommends that this zone be reset. Refer to section 5.4. 1 — **Finish Zone Recommended (FZR):** If this bit is set to '1', then the controller recommends that this zone be finished. Refer to section 5.5. 0 — **Zone Finished by Controller (ZFC):** If this bit is set to '1', then the controller finished this zone due to a Zone Active Excursion. Refer to section 5.6.
03	**Zone Attributes Information (ZAI):** Indicates additional information associated with attributes for the Zone: **Bits** / **Description** 7:4 — Reserved 3:2 — **Reset Zone Recommended Time Limit (RZRTL):** If the Reset Zone Recommended bit is set to '1', then the value in this field selects a field in the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to '1' for that zone: • A 00b value indicates the Reset Recommended Limit (RRL) field; • A 01b value indicates the Reset Recommended Limit 1 (RRL1) field; • A 10b value indicates the Reset Recommended Limit 2 (RRL2) field; and • A 11b value indicates the Reset Recommended Limit 3 (RRL3) field. 1:0 — **Finish Zone Recommended Time Limit (FZRTL):** If the Finish Zone Recommended bit is set to '1', then the value in this field selects a field in the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that indicates the amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to '1' for that zone: • A 00b value indicates the Finish Recommended Limit (FRL) field; • A 01b value indicates the Finish Recommended Limit 1 (FRL1) field; • A 10b value indicates the Finish Recommended Limit 2 (FRL2) field; and • A 11b value indicates the Finish Recommended Limit 3 (FRL3) field.

07:04	Reserved
15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).
23:16	**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.
31:24	**Write Pointer (WP):** This field is the logical block address where the next write operation for this zone should be issued. Refer to section 2.1.1.2.1 for the behavior of the write pointer.
63:32	Reserved

ZNS Specification, Section 3.4.2.2.3[111]

876. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 2 of the '772 Patent.

877. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

878. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

879. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

880. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

881. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

882. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities

[111] *Id*. at 28 (annotated).

in the United States.

883. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

884. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

885. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

886. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

887. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

888. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

889. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

890. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

891. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

892. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

893. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

894. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

895. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

896. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

897. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

898. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or

otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

899. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

900. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and

libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

901. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

902. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

903. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

904. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

905. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a

material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

906. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these

components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

907. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

908. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

909. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

910. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

911. SEC's infringement of the Asserted Patent is ongoing.

912. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

913. SEC's ongoing infringement of the Asserted Patent is willful.

914. SEC's ongoing infringement of the Asserted Patent is egregious.

915. SEA's infringement of the Asserted Patent is ongoing.

916. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

917. SEA's ongoing infringement of the Asserted Patent is willful.

918. SEA's ongoing infringement of the Asserted Patent is egregious.

919. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

920. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

921. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

922. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

923. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

924. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT THREE: INFRINGEMENT OF THE '614 PATENT

925. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

926. U.S. Patent No. 11,681,614 ("the '614 Patent"), entitled "Storage Device with Subdivisions, Subdivision Query, and Write Operations," was legally and duly issued on June 20, 2023, naming Andrey V. Kuzmin and James G. Wayda as the inventor. *See* Exhibit C.

927. Radian owns all rights, title, and interest in the '614 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future infringement.

928. The '614 Patent is valid, enforceable, and directed to patentable subject matter.

929. Radian has complied with 35 U.S.C. § 287 with respect to the '614 Patent.

930. The '614 Patent described challenges faced by conventional memory devices, including the difficulties to realize the multi-plane benefits. See, for example:

> Unfortunately, it is difficult to realize the benefits of multi-plane or multi-die architecture at an application or operating system level. That is to say, logical addresses typically arrive at the memory controller from the host as a stream of random accesses, and are sequentially assigned to first available physical space by the memory controller; as memory, particularly flash memory, is erased and recycled via read, write, wear leveling, garbage collection and other processes, sequential logical addresses become scattered throughout physical memory space. For multi-plane memory therefore, there is no practical mechanism for the host or memory controller to group related data in a manner geared for multi-plane access (i.e., there is no guarantee available physical addresses used for related data will be consistent with device multi-plane addressing restrictions).

> A need therefore exists for improvements that provide additional flexibility to systems and application designers. Ideally, such improvements would provide flexibility in storing and accessing multiple pages of data, for example, across multiple dies or planes.[112]

931. The '614 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> [W]ith insight into geometry, a host can organize address space in a manner geared for improved memory access, e.g., according to a policy aimed at IO optimization for managing multiple or linked pages.[113]

> In applying these techniques, where allocation requests allow, the host assigns physical space for writes so as to facilitate both immediate write accesses and later multiple page read access,

---

[112] Ex. C, U.S. Patent No. 11,681,614, at 2:7–41.

[113] *Id*. at 4:61–64.

> all with the benefit of the improved throughput. For example, later accesses can be performed in a manner unencumbered by memory controller address translation requirements; for multi-plane memory and shingled drives, data placement can be pre-planned in a manner consistent with any inter-plane addressing restrictions for multi-array accesses.[114]
>
> Optional features mentioned above are supported by structure where the memory controller also stores information specific to each of plural subdivisions of memory (e.g., for each physical memory unit managed by the host, in on-board registers reserved for this purpose). The memory controller makes data based on this stored information accessible to the host.[115]

932. Some of these described solutions are reflected in the claimed inventions of the '614 Patent.

933. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the '614 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the '614 Patent.

934. Claim 1 of the '614 Patent recites:

[preamble] A storage device comprising:

[1a] flash memory divided into zones, each zone of the zones being a subdivision comprising a set of erase units, and each erase unit of the set of erase units comprising host accessible pages, wherein each zone:

[1b] is associated with a logical address range,

---

[114] *Id*. at 5:1–10.

[115] *Id*. at 12:4–10.

[1c] has a zone size for which the storage device is operable to receive a management command from a host that results in the storage device performing an erasure of one or more erase units of the set of erase units,

[1d] is in a full state when there are no host accessible pages available for writing, and

[1e] is in an empty state when all host accessible pages are available for writing; and

[1f] logic operable to:

[1g] return information about the zones in response to a query command from the host, the information about the zones comprising a list of the zones accessible by the host, a zone state, the zone size, logical addresses associated with respective zones, and information representing an extent to which the respective zones include host accessible pages available for writing;

[1h] write data to the set of erase units of a specified zone using a multi-plane write command such that the data is written concurrently to two or more erase units of the set of erase units; and

[1i] update the information representing the extent to which the specified zone includes host accessible pages available for writing,

[1j] wherein the logic comprises at least one of hardware logic or instructions stored on non-transitory machine-readable media that, when executed, control operation of hardware circuitry.

935. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

936. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

937. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

938. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

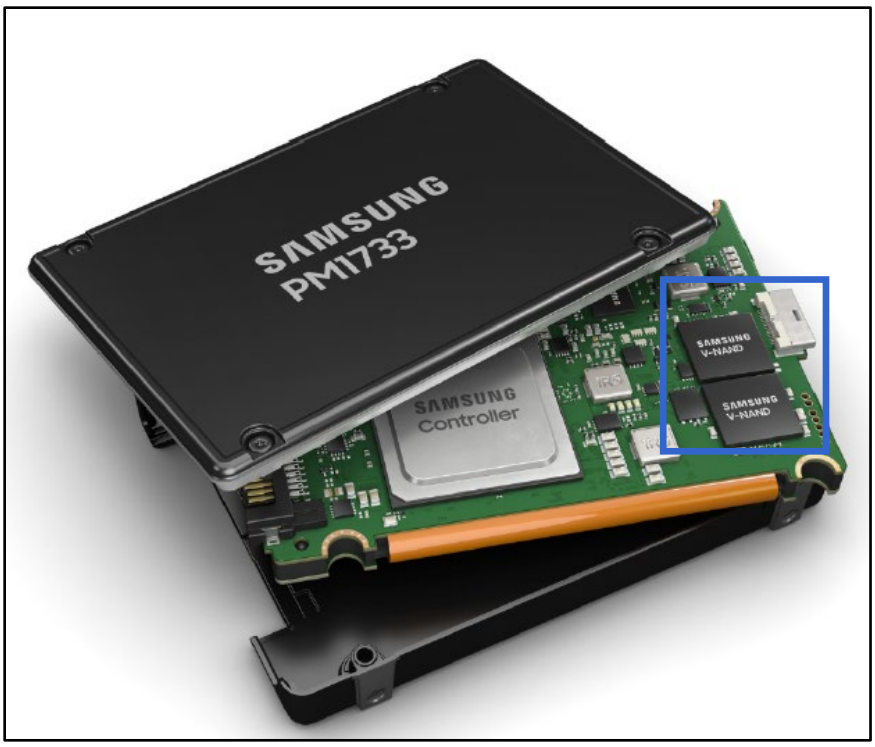
939. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

940. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

941. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a storage device. '614 Patent, claim 1, preamble.

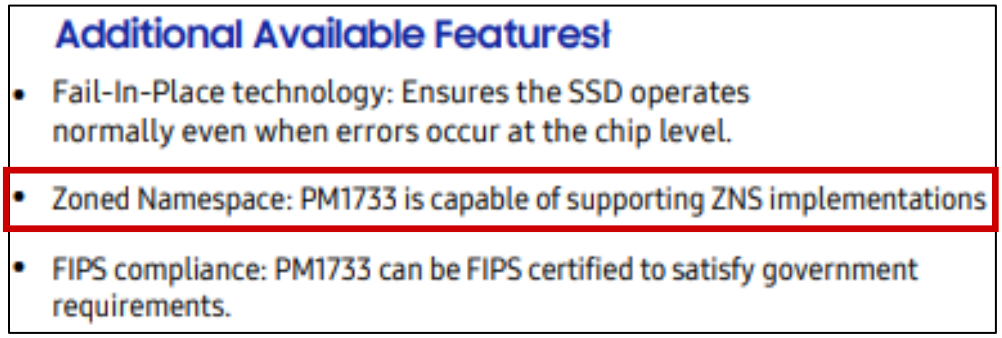
942. Samsung's Accused Instrumentalities contain flash memory divided into zones, each zone of the zones being a subdivision comprising a set of erase units, and each erase unit of the set of erase units comprising host accessible pages. '614 Patent, claim 1, element [1a].

943. For example, Samsung's PM1733 contains flash memory. See:



Samsung PM1733 Product Brief[116]

944. Samsung's Accused Instrumentalities are able to divide memory into zones. See:



Samsung PM1733 Product Brief[117]

[116] Ex. Q, PM1733 Product Brief (annotated).

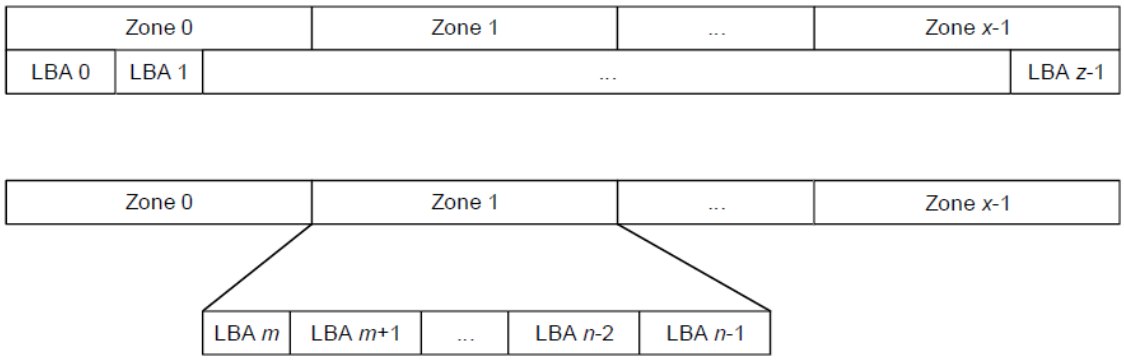
[117] *Id.* (annotated).

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



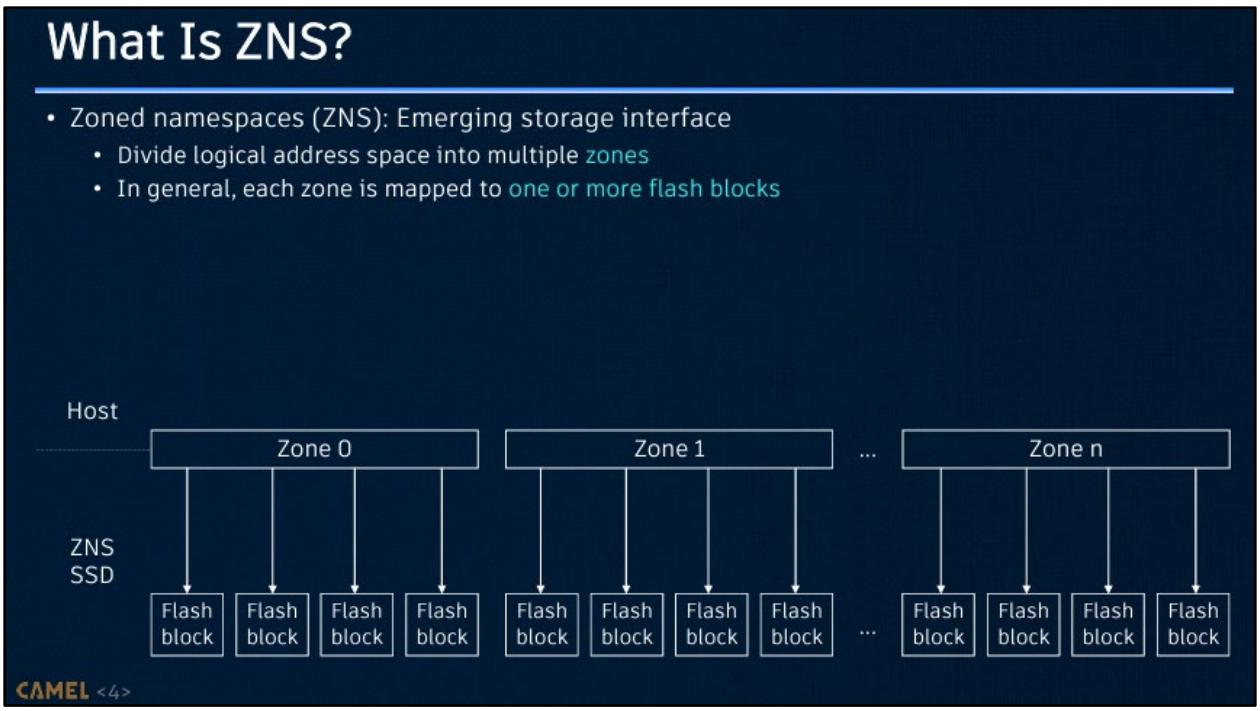
The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[118]

945. Each zone is a subdivision comprising a set of blocks (erase units). See:

[118] Ex. Z, ZNS Specification, at 8.



Presentation on ZNS[119]

> This new attribute was introduced to allow for the zone size to remain a power of two number of logical blocks (facilitating logical block to zone number conversions) while allowing optimized mapping of a zone storage capacity to the underlying media characteristics. For instance, in the case a flash based device, a zone capacity can be aligned to the size of flash erase blocks without requiring that the device implements a power-of-two sized erased block.

SSDs with NVMe Zoned Namespace (ZNS) Support[120]

946. Each erase unit comprises host-accessible pages. See:

[119] Ex. X, *What You Can't Forget: Exploiting Parallelism for Zoned Namespaces*, at 4.

[120] Ex. Y, *SSDs with NVMe Zoned Namespace (ZNS) Support*.

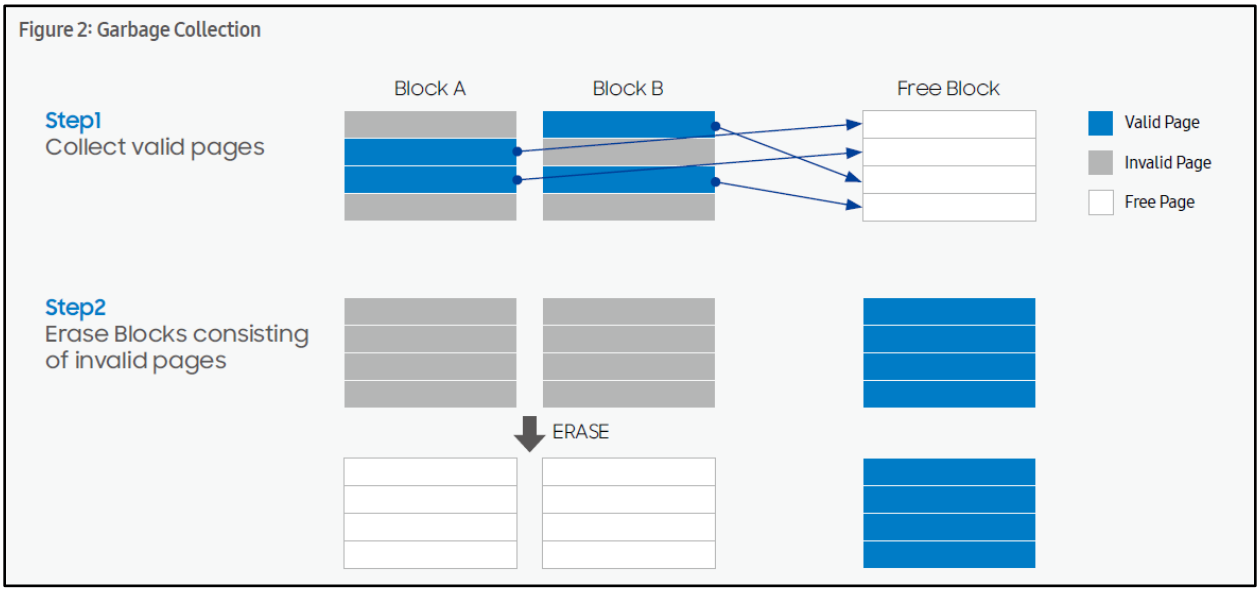
Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[121]

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Figure 2: Garbage Collection



[121] Ex. I, Samsung SSD White Paper, at 17.

Samsung Over-Provisioning White Paper[122]

2.1.3. *Simplified hierarchical architecture of a NAND flash memory chip*

Figure 2.2 illustrates a simplified architecture of a NAND flash memory chip. This is a high-level vision, relatively abstract with respect to the micro-architectural level. A flash memory chip has a hierarchical structure. It is composed of a certain number of *planes*. Planes contain *blocks*, and blocks contain *pages*. The chip is equipped with an input/output (I/O) bus which makes the reception of commands and addresses possible, as well as the transfer of data to or from the host computer system. This bus is usually multiplexed, which means that commands, addresses and data share the same pins of the flash chip. Depending on the chip version, buses with 8 or 16 bits can be found.

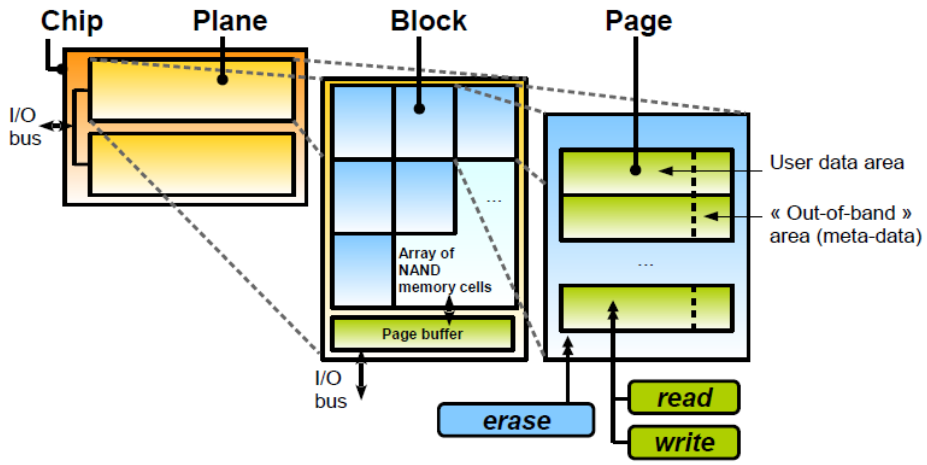


Figure 2.2. *Simplified architecture of a NAND flash memory chip. For a color version of this figure, see www.iste.co.uk/boukhobza/flash.zip*

Flash Memory Integration[123]

[122] Ex. V, Samsung Over-Provisioning White Paper, at 2.

[123] Ex. W, FLASH MEMORY INTEGRATION, at 18 (highlighted).

> – *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[124]

947. Samsung's Accused Instrumentalities contain zones wherein each zone is associated with a logical address range. '614 Patent, claim 1, element [1b].

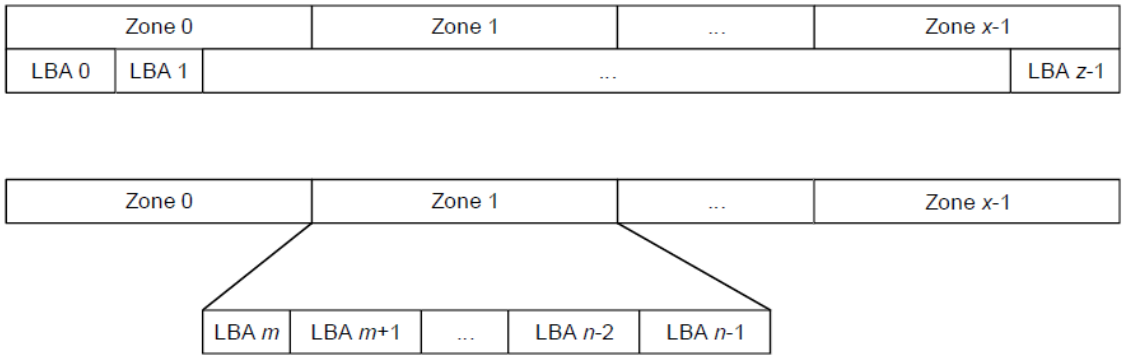
948. Each zone is associated with a range of LBAs. See:

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

[124] *Id*. at 22.

ZNS Specification, Section 2.1.1[125]

15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).
23:16	**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.

ZNS Specification, Section 3.4.2.2.3[126]

63:00	**Zone Size (ZSZE):** This field contains the size of each zone in the zoned namespace. The value is reported as a number of logical blocks. The value of the field shall not be cleared to 0h.

ZNS Specification, Section 4.1.5.1[127]

949. Samsung's Accused Instrumentalities contain zones wherein each zone has a zone size for which the storage device is operable to receive a management command from a host that results in the storage device performing an erasure of one or more erase units of the set of erase units. '614 Patent, claim 1, element [1c].

950. Each zone has a zone size. See:

[125] Ex. Z, ZNS Specification, at 8.

[126] *Id*. at 28.

[127] *Id*. at 38.

15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).
23:16	**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.

ZNS Specification, Section 3.4.2.2.3[128]

951. Samsung's Accused Instrumentalities are able to receive and execute a zone management send command specifying a reset zone action (a management command) from the host that results in the storage device performing an erasure of one or more erase units within the zone. See:

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

[128] *Id*. at 28.

ZNS Specification, Section 3.4.3[129]

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[130]

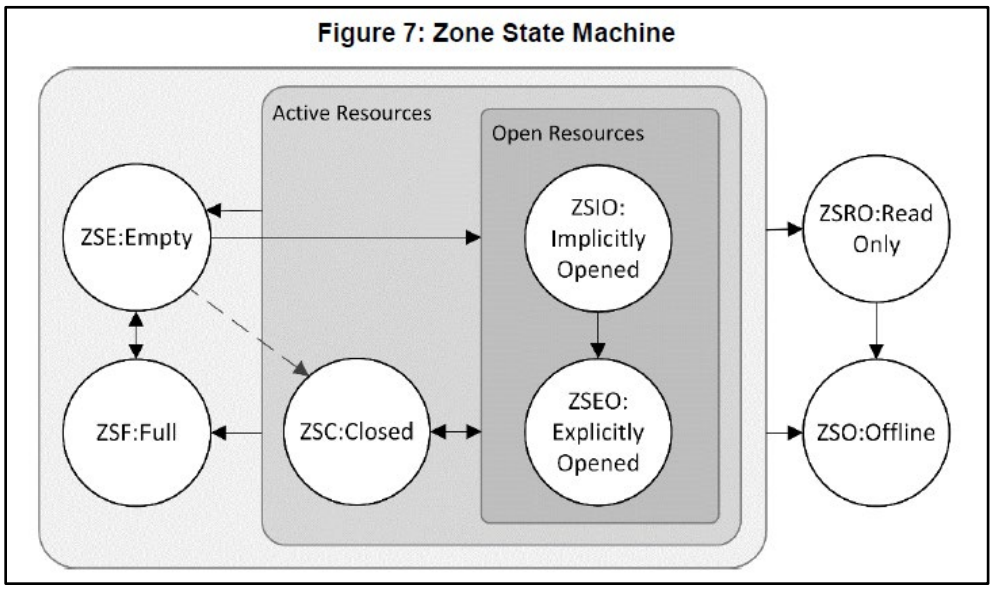
952. Samsung's Accused Instrumentalities contain zones wherein each zone is in a full state when there are no host accessible pages available for writing. '614 Patent, claim 1, element [1d].

953. A zone is in a full state when there are no host accessible pages available for

[129] *Id*. at 29.

[130] *Id*. at 30–31.

writing. See:



ZNS Specification, Section 2.1.1.3[131]

Transition ZSIO:ZSF: The zone shall transition from the ZSIO:Implicitly Opened state to the ZSF:Full state:

a) as a result of successful completion of a Zone Management Send command with a Zone Send Action of Finish Zone;
b) as a result of a write operation that writes one or more logical blocks that causes the zone to reach its writeable zone capacity;
c) due to a Zone Active Excursion (refer to section 5.6); and
d) as a result of the zoned namespace becoming write protected (refer to the Namespace Write Protection section in the NVM Express Base Specification).

ZNS Specification, Section 2.1.1.3[132]

[131] *Id*. at 12.

[132] *Id*. at 13.

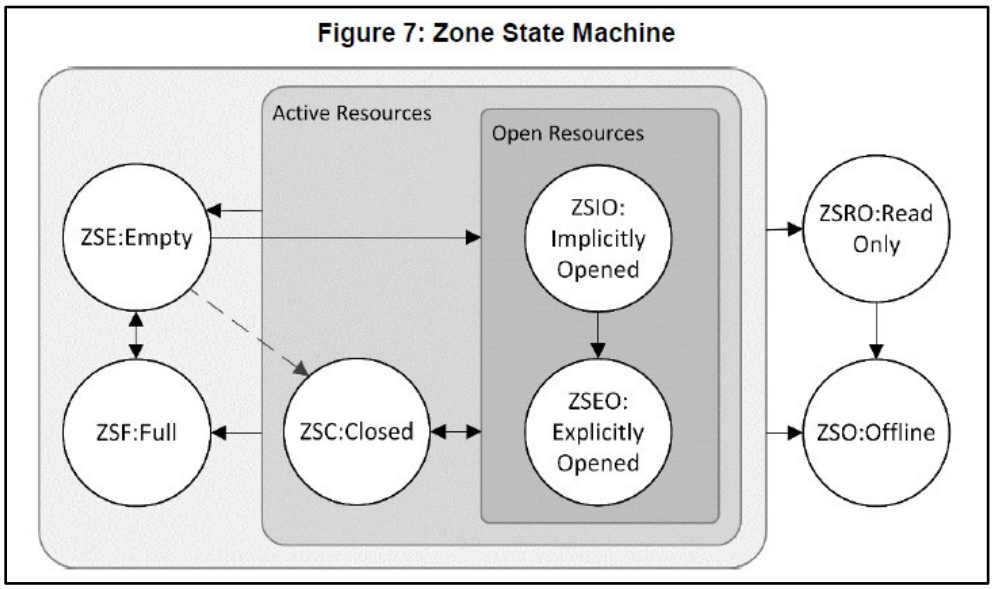
Transition ZSEO:ZSF: The zone shall transition from the ZSEO:Explicitly Opened state to the ZSF:Full state:

a) as a result of successful completion of a Zone Management Send command with a Zone Send Action of Finish Zone;
b) as a result of a write operation that writes one or more logical blocks that causes the zone to reach its writeable zone capacity;
c) due to a Zone Active Excursion (refer to section 5.6); or
d) as a result of the zoned namespace becoming write protected (refer to the Namespace Write Protection section in the NVM Express Base Specification).

ZNS Specification, Section 2.1.1.3[133]

954. Samsung's Accused Instrumentalities contain zones wherein each zone is in an empty state when all host accessible pages are available for writing. '614 Patent, claim 1, element [1e].

955. A zone is in an empty state when all host accessible pages are available for writing. See:

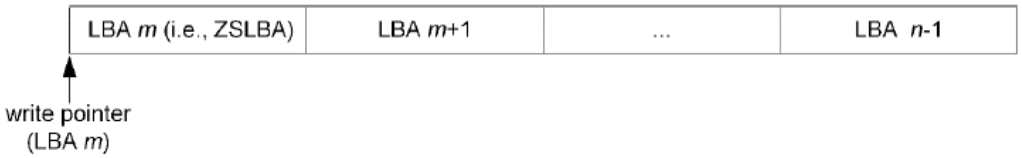


[133] *Id.*

ZNS Specification, Section 2.1.1.3[134]

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone



ZNS Specification, Section 2.1.1.2.1.1[135]

The initial state for each zone is the:

a) ZSE:Empty state, if the write pointer is valid, the write pointer points to the lowest LBA in the zone, and the Zone Descriptor Extension Valid bit is cleared to '0';

ZNS Specification, Section 2.1.1.3[136]

Transition ZSIO:ZSE: The zone shall transition from the ZSIO:Implicitly Opened state to the ZSE:Empty state as a result of successful completion of a Zone Management Send command with a Zone Send Action of Reset Zone.

ZNS Specification, Section 2.1.1.3[137]

Transition ZSEO:ZSE: The zone shall transition from the ZSEO:Explicitly Opened state to the ZSE:Empty state as a result of successful completion of a Zone Management Send command with a Zone Send Action of Reset Zone.

ZNS Specification, Section 2.1.1.3[138]

---

[134] *Id*. at 12.

[135] *Id*. at 10.

[136] *Id*. at 11.

[137] *Id*. at 12.

[138] *Id*. at 13.

> **Transition ZSC:ZSE:** The zone shall transition from the ZSC:Closed state to the ZSE:Empty state as a result of successful completion of a Zone Management Send command with a Zone Send Action of Reset Zone.

ZNS Specification, Section 2.1.1.3[139]

> **Transition ZSF:ZSE:** The zone shall transition from the ZSF:Full state to the ZSE:Empty state as a result of successful completion of a Zone Management Send command with a Zone Send Action of Reset Zone.

ZNS Specification, Section 2.1.1.3[140]

> All logical blocks in a zone shall be marked as deallocated when the zone is in the ZSE:Empty state.

ZNS Specification, Section 2.1.1.5[141]

956. Samsung's Accused Instrumentalities contain the claim element "logic operable to." '614 Patent, claim 1, element [1f]. Samsung's Accused Instrumentalities contain logic operable to perform elements [1g] through [1j]. See below.

957. Samsung's Accused Instrumentalities contain logic operable to return information about the zones in response to a query command from the host, the information about the zones comprising a list of the zones accessible by the host, a zone state, the zone size, logical addresses associated with respective zones, and information representing an extent to which the respective zones include host accessible pages available for writing. '614 Patent, claim 1, element [1g].

958. Samsung's Accused Instrumentalities are able to receive and execute a zone management receive command. See:

---

[139] *Id*.

[140] *Id*. at 14.

[141] *Id*. at 15.

> **3.4.2 Zone Management Receive command**
>
> The Zone Management Receive command returns a data buffer that contains information about zones. That information includes characteristics of the zone, the state of the zone, the capacity of the zone, and other information defined in section 3.4.2.2. The host uses this command to determine the current settings for this information.

ZNS Specification, Section 3.4.2[142]

959. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', the memory device returns information to list all zones, as well as the zone state, zone capacity (zone size), zone start LBA, and write pointer of all zones. See:

> **3.4.2 Zone Management Receive command**
>
> The Zone Management Receive command returns a data buffer that contains information about zones. That information includes characteristics of the zone, the state of the zone, the capacity of the zone, and other information defined in section 3.4.2.2. The host uses this command to determine the current settings for this information.
>
> If this information changes (e.g., as indicated by the Capacity Changed bit set to '1' in the completion queue entry of a Zone Management Send command, or by a Zone Descriptor Changed event), then the host may use this command to determine the current state of this information (e.g., the current capacity of the zone or the Reset Zone Recommended attribute).
>
> The Zone Management Receive command uses the Data Pointer, Command Dword 10, Command Dword 11, Command Dword 12, and Command Dword 13 fields. All other command specific fields are reserved.

ZNS Specification, Section 3.4.2[143]

[142] *Id*. at 24.

[143] *Id*.

15:08	**Zone Receive Action Specific:** See table below

Zone Receive Action	Description
Report Zones	**Reporting Options:** See table below
Extended Report Zones	**Reporting Options:** See table below
All other values	Reserved

Value	Description
0h	List all zones.
1h	List the zones in the ZSE:Empty state.
2h	List the zones in the ZSIO:Implicitly Opened state.
3h	List the zones in the ZSEO:Explicitly Opened state.
4h	List the zones in the ZSC:Closed state.
5h	List the zones in the ZSF:Full state.
6h	List the zones in the ZSRO:Read Only state.
7h	List the zones in the ZSO:Offline state.
8h	Reserved
9h	List all zones that have the zone attribute in the Zone Attribute field of the Zone Descriptor data structure (refer to Figure 37): • Reset Zone Recommended bit set to '1'; • Finish Zone Recommended bit set to '1'; or • Zone Finished by Controller bit set to '1'.
10h	Reserved
11h	
3Fh	
All other values	Reserved

Bits	Description	
	00h	**Report Zones:** Reports Zone Descriptor entries through the Report Zones data structure (refer to Figure 35).
	01h	**Extended Report Zones:** Reports Zone Descriptor entries through the Extended Report Zones data structure (refer to Figure 36). This value is supported if the zoned namespace is formatted with a non-zero Zone Descriptor Extension Size. Otherwise, the controller shall abort the command with a status code of Invalid Field in Command.
	02h to FFh	Reserved

ZNS Specification, Section 3.4.2[144]

[144] *Id*. at 25–26.

Figure 35: Report Zones Data Structure

Bytes	Description
07:00	**Number of Zones:** If the Partial Report bit (refer to Figure 34) is cleared to '0', then this field indicates the number of zones that match the criteria defined in section 3.4.2.1.1. If the Partial Report bit is set to '1', then this field indicates the number of zones for which complete Zone Descriptors were transferred to the data buffer. Refer to section 3.4.2.1.1 for the content of the data buffer.
63:08	Reserved
127:64	**Zone Descriptor 0:** Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).
191:128	**Zone Descriptor 1:** Contains the Zone Descriptor for the second zone reported, if any.
…	…
((n+1)*64)+63:(n+1)*64	**Zone Descriptor n:** Contains the Zone Descriptor for the last zone reported, if any.

ZNS Specification, Section 3.4.2.2.1[145]

3.4.2.2.3 Zone Descriptor Data Structure

Figure 37 defines the Zone Descriptor data structure.

Figure 37: Zone Descriptor Data Structure

Bytes	Description	
00	Bits	Description
	7:4	Reserved
	3:0	**Zone Type (ZT):** This field indicates the type of the zone. Value: 2h — Definition: Sequential Write Required Value: All other values — Definition: Reserved
01	Bits	Description
	7:4	**Zone State (ZS):** This field indicates the state of the zone. Value: 0h — Definition: Reserved Value: 1h — Definition: Empty Value: 2h — Definition: Implicitly Opened Value: 3h — Definition: Explicitly Opened Value: 4h — Definition: Closed Value: 5h to Ch — Definition: Reserved Value: Dh — Definition: Read Only Value: Eh — Definition: Full Value: Fh — Definition: Offline
	3:0	Reserved

[145] *Id*. at 26.

Figure 37: Zone Descriptor Data Structure

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>02</td><td>Zone Attributes (ZA): Indicates attributes for the Zone:
<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7</td><td>Zone Descriptor Extension Valid (ZDEV): If this bit is set to ‘1’, then Zone Descriptor Extension data is associated with the zone. If this bit is cleared to ‘0’, then no Zone Descriptor Extension data is associated with the zone. Refer to section 5.1.</td></tr><tr><td>6:3</td><td>Reserved</td></tr><tr><td>2</td><td>Reset Zone Recommended (RZR): If this bit is set to ‘1’, then the controller recommends that this zone be reset. Refer to section 5.4.</td></tr><tr><td>1</td><td>Finish Zone Recommended (FZR): If this bit is set to ‘1’, then the controller recommends that this zone be finished. Refer to section 5.5.</td></tr><tr><td>0</td><td>Zone Finished by Controller (ZFC): If this bit is set to ‘1’, then the controller finished this zone due to a Zone Active Excursion. Refer to section 5.6.</td></tr></table></td></tr>
<tr><td>03</td><td>Zone Attributes Information (ZAI): Indicates additional information associated with attributes for the Zone:
<table><tr><th>Bits</th><th>Description</th></tr><tr><td>7:4</td><td>Reserved</td></tr><tr><td>3:2</td><td>Reset Zone Recommended Time Limit (RZRTL): If the Reset Zone Recommended bit is set to ‘1’, then the value in this field selects a field in the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that indicates amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Reset Zone Recommended zone attribute is set to ‘1’ for that zone:
A 00b value indicates the Reset Recommended Limit (RRL) field;
A 01b value indicates the Reset Recommended Limit 1 (RRL1) field;
A 10b value indicates the Reset Recommended Limit 2 (RRL2) field; and
A 11b value indicates the Reset Recommended Limit 3 (RRL3) field.</td></tr><tr><td>1:0</td><td>Finish Zone Recommended Time Limit (FZRTL): If the Finish Zone Recommended bit is set to ‘1’, then the value in this field selects a field in the I/O Command Set specific Identify Namespace data structure for the Zoned Namespace Command Set that indicates the amount of time before the NVM subsystem may perform a vendor specific action on a zone after the Finish Zone Recommended zone attribute is set to ‘1’ for that zone:
A 00b value indicates the Finish Recommended Limit (FRL) field;
A 01b value indicates the Finish Recommended Limit 1 (FRL1) field;
A 10b value indicates the Finish Recommended Limit 2 (FRL2) field; and
A 11b value indicates the Finish Recommended Limit 3 (FRL3) field.</td></tr></table></td></tr>
<tr><td>07:04</td><td>Reserved</td></tr>
<tr><td>15:08</td><td>Zone Capacity (ZCAP): This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49).

If the Variable Zone Capacity bit is cleared to ‘0’ in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace.

If the Variable Zone Capacity bit is set to ‘1’ in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).</td></tr>
<tr><td>23:16</td><td>Zone Start Logical Block Address (ZSLBA): This field contains the 64-bit address of the lowest logical block for the zone.</td></tr>
<tr><td>31:24</td><td>Write Pointer (WP): This field is the logical block address where the next write operation for this zone should be issued. Refer to section 2.1.1.2.1 for the behavior of the write pointer.</td></tr>
<tr><td>63:32</td><td>Reserved</td></tr>
</table>

ZNS Specification, Section 3.4.2.2.3[146]

960. The write pointer is information representing an extent to which a particular zone includes host accessible pages available for writing. See:

> **2.1.1.2.1.1 Writing in Sequential Write Required Zones**
>
> The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:
>
> - Write command;
> - Write Zeroes command;
> - Write Uncorrectable command;
> - Copy command; and
> - Zone Append command.
>
> A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

ZNS Specification, Section 2.1.1.2.1.1[147]

[146] *Id*. at 27–28.

[147] *Id*. at 9.

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

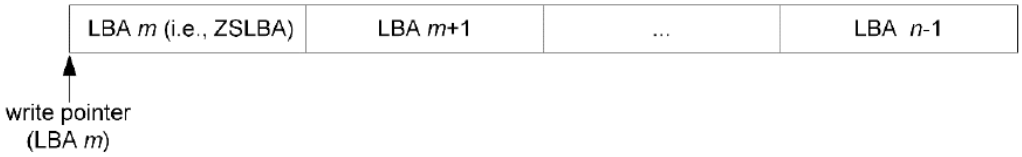
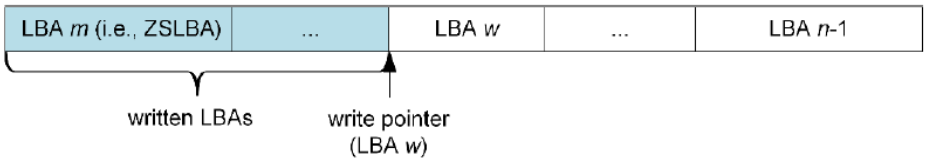


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[148]

961. Samsung's Accused Instrumentalities contain logic operable to write data to the set of erase units of a specified zone using a multi-plane write command such that the data is written concurrently to two or more erase units of the set of erase units. '614 Patent, claim 1, element [1h].

962. Samsung's Accused Instrumentalities are able to receive and execute a write command from the host. See:

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

[148] *Id*. at 10.

ZNS Specification, Section 3.3.7[149]

963. Samsung's Accused Instrumentalities are also able to receive and execute a Zone Append command from the host. See:

> **3.4.1 Zone Append command**
>
> The Zone Append command writes data and metadata, if applicable, to the I/O controller for the zone indicated by the ZSLBA field. The controller assigns the data and metadata, if applicable, to a set of logical blocks within the zone. The lowest LBA of the set of logical blocks written is returned in the completion queue entry (refer to section 3.4.1.2). The host may also specify protection information to include as part of the operation.
>
> This command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used. All other command specific fields are reserved.

ZNS Specification, Section 3.3.7[150]

964. Samsung's Accused Instrumentalities perform a multi-plane write such that the data is written concurrently to two or more erase units of a zone. See:

> **Memory density and die floor plan**
>
> Figure 1 shows 16 48L V-NAND dice with two F-Chips in an MCP (multichip package). Die efficiency is higher in the 48L. The 32L V-NAND die area is 84.3 mm^2 while the 48L V-NAND die measures 99.8 mm^2, which is a 17.3% increase due to greater die length (Figure 2). Memory density per unit die area is increased to 2.57 Gb/mm^2. The memory density of leading-edge 2D planar NAND devices such as Toshiba 15 nm TLC NAND is 1.28 Gb/mm^2, (see more on our analysis on the Toshiba 15 nm NAND). The key feature differences in die floor plans are 1) planar (NAND memory array) area, 2) bitline switch and page buffer area, 3) LOGIC and peripheral area, and 4) the addition of F-chip. Each die has two planes. The NAND memory array area increased from 48.9 mm^2 to 68.7 mm^2 which is 40.3% larger. The area of the bitline switch circuit is the same as that of the 32L, while the page buffer area is 20% smaller. The logic and peripheral circuit area is 34.8% reduced. In other words, Samsung dramatically shrank the area of the page buffer and peripheral region so that they can further increase memory density and die efficiency. Die thickness is also reduced from 132 μm to 36 μm for the sixteen die stacked in the MCP.

[149] *Id*. at 20.

[150] *Id*. at 21.

EETimes V-NAND analysis[151]

Samsung TLC V- NAND Flash Memory		
	9th Gen V-NAND	8th Gen V-NAND
Layers	290?	236
Decks	2 (x145)	2 (x118)
Die Capacity	1 Tbit	1 Tbit
Die Size (mm^2)	?mm2	?mm2
Density ($Gbit/mm^2$)	?	?
I/O Speed	3.2 GT/s (Toggle 5.1)	2.4 GT/s (Toggle 5.0)
Planes	6?	4
CuA / PuC	Yes	Yes

AnandTech V-NAND analysis [152]

To maximize the flash operation parallelism, the flash blocks from a set of flash chips accessible in parallel will compose the FBG of a zone, and the chunks of a zone need to be interleavingly placed on the parallel flash chips. The number of parallel flash chips for the chunk interleaving is referred to as the *zone interleaving degree* D_{zone}, and the set of logically consecutive chunks across the parallel flash chips is referred to as a **stripe**. For a coarse-grained zone-to-FBG mapping, an FBG has flash blocks at the same block offset in the parallel flash chips, and the chunks of a stripe are located at the same page offset in different flash blocks.

Samsung Research Paper[153]

---

[151] Ex. FF, *Samsung's 3D V-NAND 32L vs 48L–Just Vertical Expansion?*

[152] Ex. GG, *Samsung Starts Mass Production of 9th Generation V-NAND: 1Tb 3D TLC NAND.*

[153] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

965. Samsung's Accused Instrumentalities contain logic operable to update the information representing the extent to which the specified zone includes host accessible pages available for writing. '614 Patent, claim 1, element [1i].

966. Samsung's Accused Instrumentalities are able to update the write pointer. See:

> **2.1.1.2.1.1 Writing in Sequential Write Required Zones**
>
> The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:
>
> - Write command;
> - Write Zeroes command;
> - Write Uncorrectable command;
> - Copy command; and
> - Zone Append command.
>
> A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

ZNS Specification, Section 2.1.1.2.1.1[154]

> The write pointer for a zone in the ZSE:Empty state, the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state shall be increased by the number of logical blocks written on successful completion of a write operation.

ZNS Specification, Section 2.1.1.2.1.1[155]

[154] Ex. Z, ZNS Specification, at 9.

[155] *Id*. at 10.

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

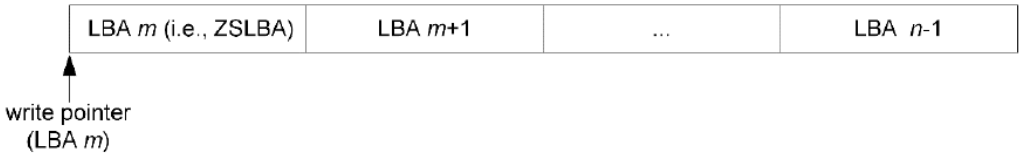
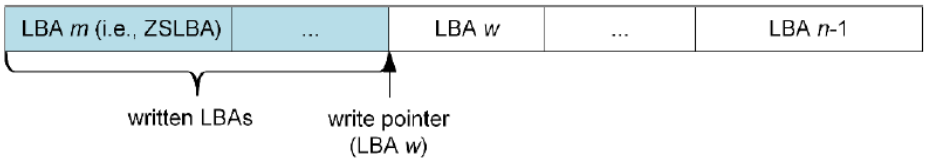


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[156]

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

[156] *Id*. at 10.

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

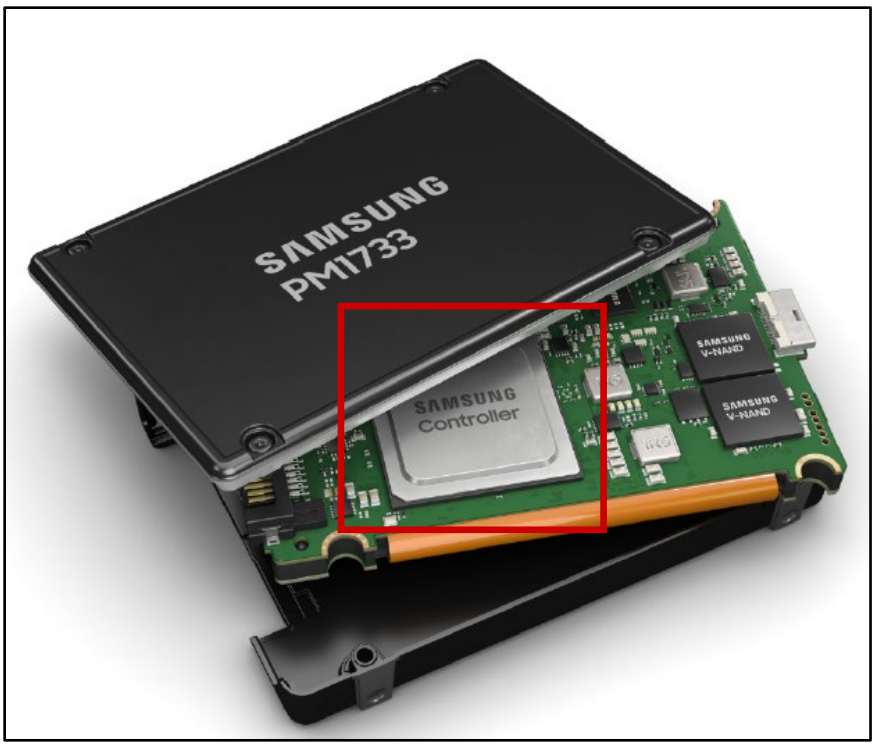
a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[157]

967. On information and belief, in Samsung's Accused Instrumentalities, said logic comprises at least one of hardware logic or instructions stored on non-transitory machine-readable media that, when executed, control operation of hardware circuitry. '614 Patent, claim 1, element [1j]. See, for example:

[157] *Id*. at 30–31.



Samsung PM1733 Product Brief[158]

968. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '614 Patent.

969. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

970. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

971. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

972. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

---

[158] Ex. Q, Samsung PM1733 Product Brief (annotated).

973. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

974. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

975. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

976. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

977. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

978. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

979. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

980. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

981. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

982. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

983. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

984. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities

in the United States.

985. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

986. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

987. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

988. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

989. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

990. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or

hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

991. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

992. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a

material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

993. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

994. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

995. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

996. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of

the Accused Instrumentalities in the United States.

997. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

998. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their

components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

999. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1000. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1001. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

1002. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

1003. SEC's infringement of the Asserted Patent is ongoing.

1004. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1005. SEC's ongoing infringement of the Asserted Patent is willful.

1006. SEC's ongoing infringement of the Asserted Patent is egregious.

1007. SEA's infringement of the Asserted Patent is ongoing.

1008. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1009. SEA's ongoing infringement of the Asserted Patent is willful.

1010. SEA's ongoing infringement of the Asserted Patent is egregious.

1011. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

1012. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

1013. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

1014. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

1015. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

1016. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT FOUR: INFRINGEMENT OF THE '801 PATENT

1017. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

1018. U.S. Patent No. 11,740,801 ("the '801 Patent"), entitled "Cooperative Flash Management of Storage Device Subdivisions," was legally and duly issued on August 29, 2023, naming Andrey V. Kuzmin, Alan Chen, and Robert Lercari as the inventor. *See* Exhibit D.

1019. Radian owns all rights, title, and interest in the '801 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future

infringement.

1020. The ’801 Patent is valid, enforceable, and directed to patentable subject matter.

1021. Radian has complied with 35 U.S.C. § 287 with respect to the ’801 Patent.

1022. The ’801 Patent described challenges faced by conventional memory devices involving data storage and migration. See, for example:

> What is needed are techniques for more efficiently managing operation of nonvolatile memory. Still more particularly, what is needed are techniques for more efficiently managing memory in a heterogeneous memory system. The present invention addresses these needs and provides further, related advantages.[159]

> This disclosure provides techniques for data storage and/or migration based on metrics generated or tracked for respective data (“per-data metrics”) and based on memory degradation.[160]

1023. The ’801 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> The system or one of its components generates information that identifies hot/cold status, age, or another metric characterizing respective data, and it also tracks life cycle wear or degradation for at least one tier of NV memory; the system (or a host or a memory controller or both) uses both these persistence metrics and/or wear in managing a memory operation.
>
> In a specific embodiment, if a minimum free space is to be preserved, a memory controller, the host or both executes a routine where data is identified in the memory that violates a hot/cold, persistence or age criterion for that memory, or alternatively, that the host, memory controller or both, deems more suitable to a different location within the same tier of memory or a different

[159] Ex. D, U.S. Patent No. 11,740,801, at 1:66–2:4.

[160] *Id*. at 4:61–64.

> memory tier; the identified data is then moved (and can be written into a different tier, as appropriate).[161]
>
> Generally speaking, a cooperative memory controller for NV memory stores information specific to each of plural subdivisions of memory, and makes data based on that stored information accessible to the host to assist with host management of memory, as needed.[162]

1024. Some of these described solutions are reflected in the claimed inventions of the ’801 Patent.

1025. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the ’801 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the ’801 Patent.

1026. Claim 1 of the ’801 Patent recites:

> [preamble] A storage device comprising:
>
> [1a] NAND flash memory; and
>
> [1b] a memory controller having logic to cause the memory controller to
>
> [1c] map non-overlapping ranges of logical addresses to respective subdivisions of physical storage in the NAND flash memory, each of the respective subdivisions comprising erase blocks that must each be physically erased as a unit,
>
> [1d] receive write requests, accompanying addresses and

[161] *Id*. at 3:23–46.

[162] *Id*. at 23:52–56.

accompanying data from a host and, with respect to each one of the write requests,

[1e] identify a specific one of the respective subdivisions dependent on which one of the non-overlapping ranges encompasses the address accompanying the one of the write requests,

[1f] program the data accompanying with the one of the write requests into a next available storage location within the specific one of the respective subdivisions, and

[1g] in association with the programming of the data accompanying the one of the write requests, update information, for the specific one of the respective subdivisions, representing an extent to which the erase blocks of the specific one of the respective subdivisions are full,

[1h] store a value representing a size of the NAND flash memory that is to be physically erased in fulfilling a maintenance request issued by the host,

[1i] provide to the host, responsive to a query received from the host, the information, as updated, and the stored value,

[1j] update metadata for the specific one of the respective subdivisions, wherein the metadata represents a characteristic of the data in the specific one of the respective subdivisions,

[1k] with respect to the metadata, as updated,

[1l] perform a comparison based on the metadata to detect a condition, and

[1m] store information, responsive to detection of the condition, wherein the information is accessible to the host and indicates that the condition was detected;

[1n] wherein each said logic comprises at least one of circuitry or instructions stored on physical storage media that, when executed, are to control operation of circuitry.

1027. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

1028. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1029. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1030. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

1031. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

1032. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

1033. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a storage device. '801 Patent, claim 1, preamble.

1034. Samsung's Accused Instrumentalities contain a NAND flash memory. '801 Patent, claim 1, element [1a].

1035. For example, Samsung's PM1733 contains NAND flash memory. See:

Samsung PM1733 specifications

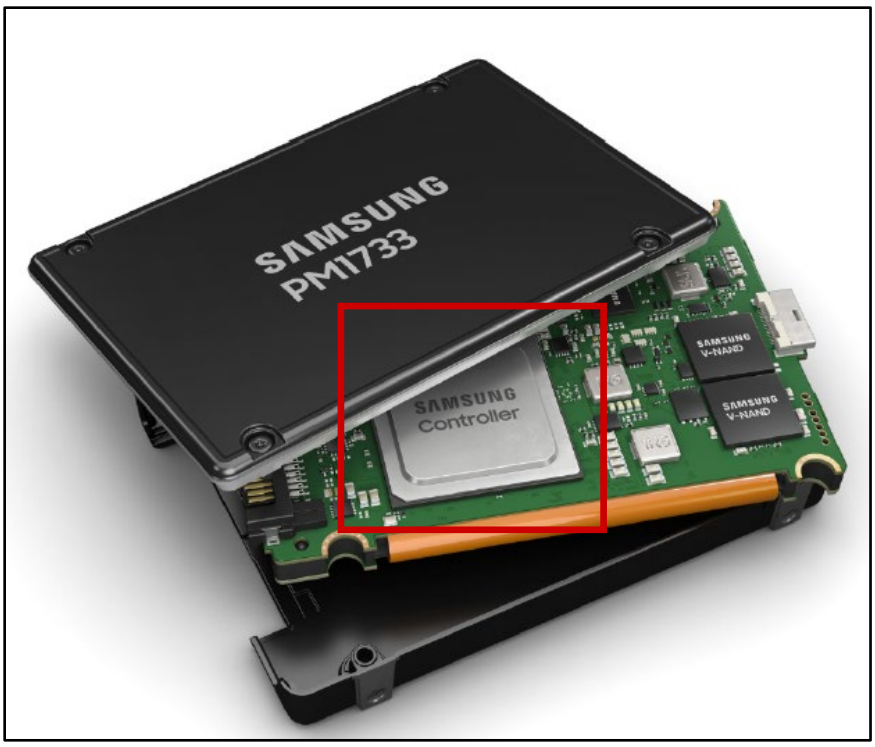
Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[163]

1036. Samsung's Accused Instrumentalities contain a memory controller having logic to cause the memory controller to perform elements [1c] through [1m]. '801 Patent, claim 1, element [1b]. See below.

1037. For example, Samsung's PM1733 contains a memory controller. See:

[163] Ex. Q, Samsung PM1733 Product Brief.



Samsung PM1733 Product Brief[164]

1038. Samsung's Accused Instrumentalities contain logic to cause the memory controller to map non-overlapping ranges of logical addresses to respective subdivisions of physical storage in the NAND flash memory, each of the respective subdivisions comprising erase blocks that must each be physically erased as a unit. '801 Patent, claim 1, element [1c].

1039. Samsung's Accused Instrumentalities contain erase blocks that must each be physically erased as a unit. See:

[164] *Id*. (annotated).

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Samsung Over-Provisioning White Paper[165]

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[166]

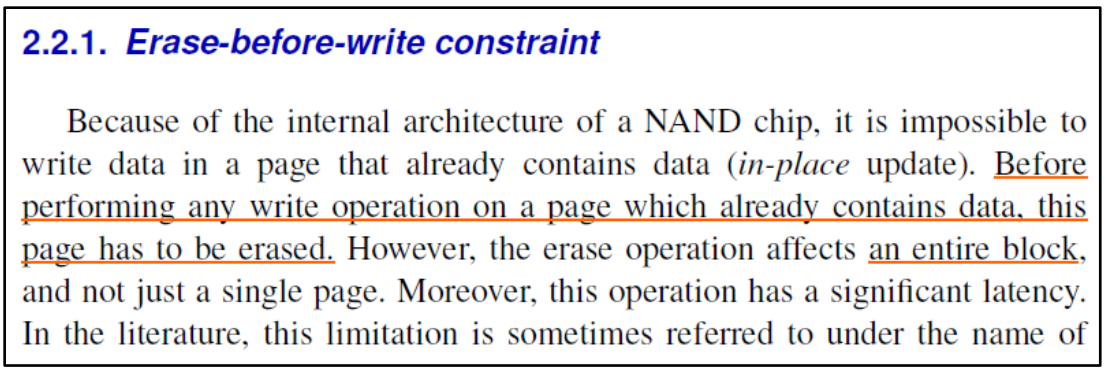
– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[167]

[165] Ex. V, Samsung Over-Provisioning White Paper, at 2 (highlighted).

[166] Ex. I, Samsung SSD White Paper, at 17.

[167] Ex. W, FLASH MEMORY INTEGRATION, at 22 (highlighted).

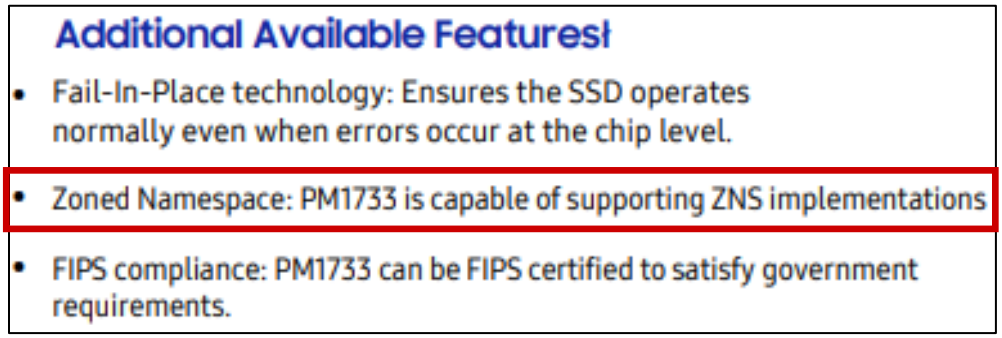


2.2.1. *Erase-before-write constraint*

Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[168]

1040. Samsung's Accused Instrumentalities are able to map non-overlapping ranges of logical addresses to respective zones (subdivisions) of physical storage in the NAND flash memory, each of the respective subdivisions comprising erase blocks. See:



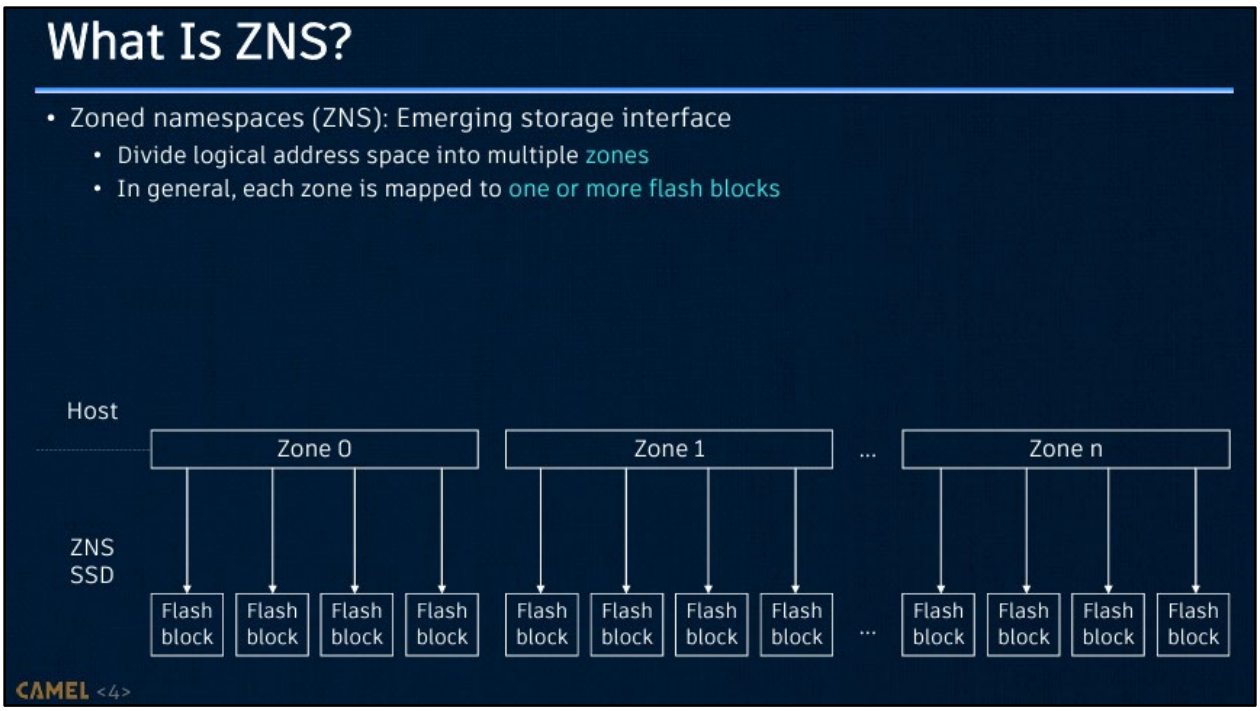
Additional Available Features!

- Fail-In-Place technology: Ensures the SSD operates normally even when errors occur at the chip level.
- Zoned Namespace: PM1733 is capable of supporting ZNS implementations
- FIPS compliance: PM1733 can be FIPS certified to satisfy government requirements.

Samsung PM1733 Product Brief[169]

[168] *Id*. at 24.

[169] Ex. Q, PM1733 Product Brief (annotated).



Presentation on ZNS[170]

> This new attribute was introduced to allow for the zone size to remain a power of two number of logical blocks (facilitating logical block to zone number conversions) while allowing optimized mapping of a zone storage capacity to the underlying media characteristics. For instance, in the case a flash based device, a zone capacity can be aligned to the size of flash erase blocks without requiring that the device implements a power-of-two sized erased block.

SSDs with NVMe Zoned Namespace (ZNS) Support[171]

[170] Ex. X, *What You Can't Forget: Exploiting Parallelism for Zoned Namespaces*, at 4.

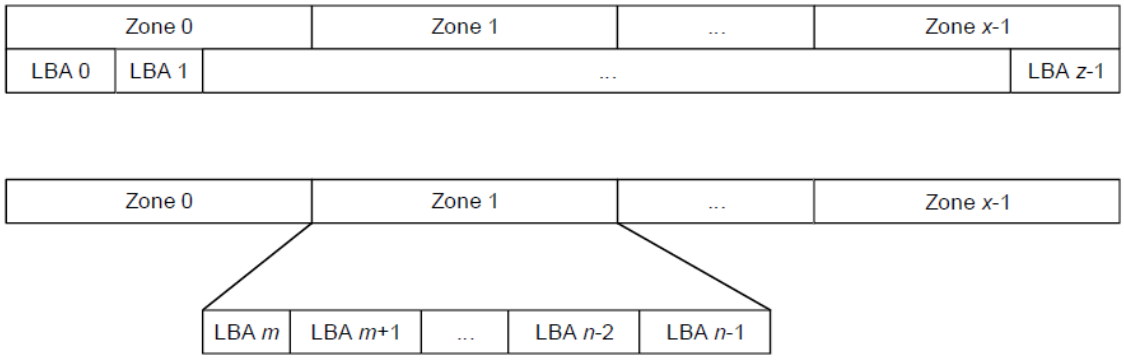
[171] Ex. Y, *SSDs with NVMe Zoned Namespace Support*.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[172]

1041. Samsung's Accused Instrumentalities contain logic to cause the memory controller to receive write requests, accompanying addresses and accompanying data from a host. '801 Patent, claim 1, element [1d].

1042. Samsung's Accused Instrumentalities are able to receive and execute Write commands from a host.

1043. A Write command is a data write request.

1044. A Write command writes data and metadata, if applicable, to the I/O controller for

[172] Ex. Z, ZNS Specification, at 8.

the logical blocks indicated. See:

> **3.3.6 Write command**
>
> The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

NVM Command Set Specification, Section 3.3.6[173]

1045. The Write command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:

[173] Ex. AA, NVM Command Set Specification, at 53.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[174]

[174] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[175]

1046. The NSID is incoming address information.

1047. The Write command is accompanied by a logical block address ("LBA") and data.

See:

[175] Ex. BB, NVMe Base Specification, at 107.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2

ZNS Specification, Section 3.2[176]

3.3.6 Write command

The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used.

Figure 67: Write – Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where data is transferred from. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

[176] Ex. Z, ZNS Specification, at 18 (annotated).

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[177]

Figure 88: Common Command Format

<table>
<tr><th>Bytes</th><th>Description</th></tr>
<tr><td>39:24</td><td>
Data Pointer (DPTR): This field specifies the data used in the command.
If CDW0.PSDT is cleared to 00b, then the definition of this field is:
<table>
<tr><td>39:32</td><td>PRP Entry 2 (PRP2): This field:
• is reserved if the data transfer does not cross a memory page boundary;
• specifies the Page Base Address of the second memory page if the data transfer crosses exactly one memory page boundary. E.g.,:
  ○ the command data transfer length is equal in size to one memory page and the offset portion of the PBAO field of PRP1 is non-zero; or
  ○ the Offset portion of the PBAO field of PRP1 is equal to 0h and the command data transfer length is greater than one memory page and less than or equal to two memory pages in size;
and
• is a PRP List pointer if the data transfer crosses more than one memory page boundary. E.g.,:
  ○ the command data transfer length is greater than or equal to two memory pages in size but the offset portion of the PBAO field of PRP1 is non-zero; or
  ○ the command data transfer length is equal in size to more than two memory pages and the Offset portion of the PBAO field of PRP1 is equal to 0h.</td></tr>
<tr><td>31:24</td><td>PRP Entry 1 (PRP1): This field contains the first PRP entry for the command or a PRP List pointer depending on the command.</td></tr>
</table>
If CDW0.PSDT is set to 01b or 10b, then the definition of this field is:
<table>
<tr><td>39:24</td><td>SGL Entry 1 (SGL1): This field contains the first SGL segment for the command. If the SGL segment is an SGL Data Block or Keyed SGL Data Block or Transport SGL Data Block descriptor, then it describes the entire data transfer. If more than one SGL segment is needed to describe the data transfer, then the first SGL segment is a Segment, or Last Segment descriptor. Refer to section 4.1.2 for the definition of SGL segments and descriptor types.
The NVMe Transport may support a subset of SGL Descriptor types and features as defined in the NVMe Transport binding specification.</td></tr>
</table>
</td></tr>
</table>

[177] Ex. AA, NVM Command Set Specification, at 53–54.

NVMe Base Specification, Section 3.3.3.1[178]

1048. The starting LBA is incoming address information.

1049. Samsung's Accused Instrumentalities are also able to receive and execute a Zone Append command from a host.

1050. A Zone Append command is a write request.

1051. The Zone Append command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:

[178] Ex. BB, NVM Express Base Specification, at 107.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[179]

[179] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[180]

1052. The Zone Append command is accompanied by data and a Zone Start Logical Block Address ("ZSLBA") that designates a specific one of the zones. See:

3.4 Zoned Namespace Command Set I/O Commands

3.4.1 Zone Append command

The Zone Append command writes data and metadata, if applicable, to the I/O controller for the zone indicated by the ZSLBA field. The controller assigns the data and metadata, if applicable, to a set of logical blocks within the zone. The lowest LBA of the set of logical blocks written is returned in the completion queue entry (refer to section 3.4.1.2). The host may also specify protection information to include as part of the operation.

This command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used. All other command specific fields are reserved.

[180] Ex. BB, NVMe Base Specification, at 107.

Figure 23: Zone Append – Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where data is transferred from. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

Figure 25: Zone Append – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Zone Start Logical Block Address (ZSLBA):** This field indicates the 64-bit address of the lowest logical block of the zone in which the data and metadata, if applicable, associated with this command is to be stored. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

ZNS Specification, Section 3.4.1[181]

1053. Samsung's Accused Instrumentalities contain logic to, with respect to each one of the write requests, identify a specific one of the respective subdivisions dependent on which one of the non-overlapping ranges encompasses the address accompanying the one of the write requests. '801 Patent, claim 1, element [1e].

1054. Samsung's Accused Instrumentalities are able to identify a specific zone dependent on which one of the zones encompasses the received starting LBA accompanying a Write command or the received zone start LBA accompanying a Zone Append command. See:

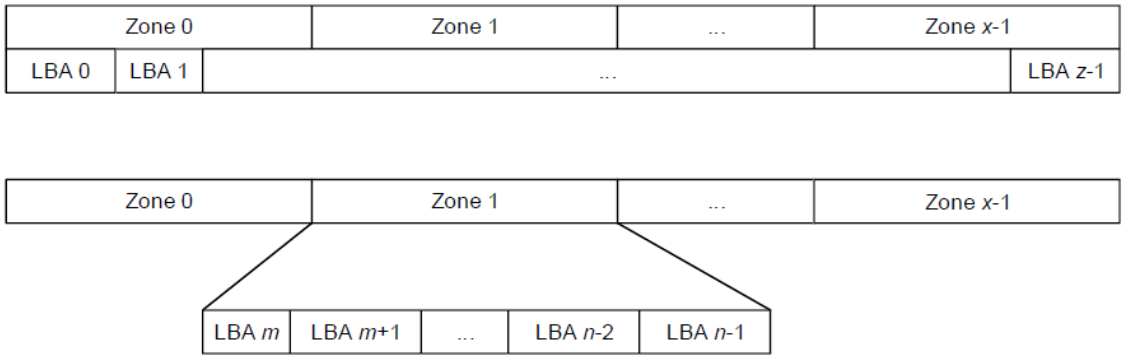
[181] Ex. Z, ZNS Specification, at 21–22.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[182]

[182] *Id.* at 8.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[183]

1055. Samsung's Accused Instrumentalities contain logic to, with respect to each one of the write requests, program the data accompanying with the one of the write requests into a next available storage location within the specific one of the respective subdivisions. '801 Patent, claim 1, element [1f].

1056. The data is to be programmed into a zone at the zone's write pointer (the next available storage location within the zone). See:

[183] *Id*. at 20.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[184]

Write Pointer	The Write Pointer attribute defines the lowest numbered writeable logical block address in that zone. The validity of the write pointer is zone state specific and is defined per zone type (refer to section 2.1.1.2).

ZNS Specification, Section 2.1.1.1[185]

2.1.1.2.1.1 Writing in Sequential Write Required Zones

The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:

- Write command;
- Write Zeroes command;
- Write Uncorrectable command;
- Copy command; and
- Zone Append command.

A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

[184] *Id*. at 20.

[185] *Id*. at 9.

ZNS Specification, Section 2.1.1.2.1.1[186]

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

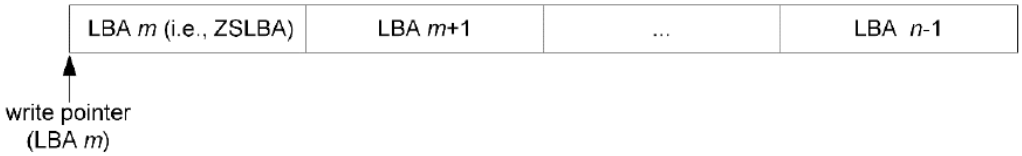
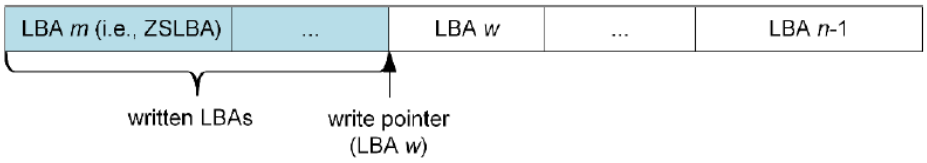


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[187]

1057. Samsung's Accused Instrumentalities contain logic to, with respect to each one of the write requests, in association with the programming of the data accompanying the one of the write requests, update information, for the specific one of the respective subdivisions, representing an extent to which the erase blocks of the specific one of the respective subdivisions are full. '801 Patent, claim 1, element [1g].

1058. Samsung's Accused Instrumentalities are able to update the write pointer in association with the programming of write data. The write pointer represents an extent to which

---

[186] *Id.*

[187] *Id.* at 10.

the erase blocks of a specific zone are full. See:

> **2.1.1.2.1.1 Writing in Sequential Write Required Zones**
>
> The following commands may be used to write to logical blocks in a specific zone of zone type Sequential Write Required:
>
> - Write command;
> - Write Zeroes command;
> - Write Uncorrectable command;
> - Copy command; and
> - Zone Append command.
>
> A write pointer is maintained for each zone in the zoned namespace that indicates the next writeable logical block address in that zone. The write pointer is valid for a subset of the zone states as defined in Figure 4.

ZNS Specification, Section 2.1.1.2.1.1[188]

> The write pointer for a zone in the ZSE:Empty state, the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state shall be increased by the number of logical blocks written on successful completion of a write operation.

ZNS Specification, Section 2.1.1.2.1.1[189]

[188] *Id*. at 9.

[189] *Id*. at 10.

Figure 5 shows an example of a zone in the ZSE:Empty state. LBA m is the ZSLBA attribute, the write pointer indicates ZSLBA, and n-1 is the highest LBA of the zone.

Figure 5: Write Pointer in an Empty Zone

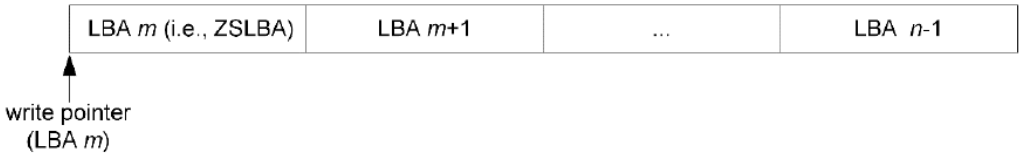
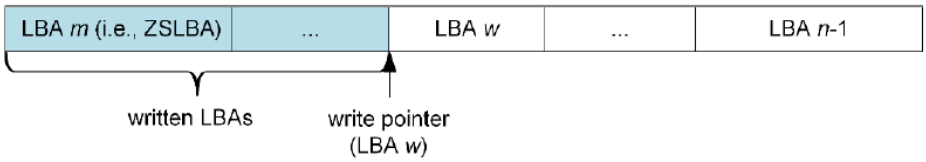


Figure 6 shows an example of a zone in the ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, or the ZSC:Closed state, that has had some logical blocks written. The write pointer, indicated by LBA *w*, is the lowest-numbered unwritten LBA (i.e., the next LBA to be written) and *n-1* is the highest LBA of the zone.

Figure 6: Write Pointer in a Partially Written Zone



ZNS Specification, Section 2.1.1.2.1.1[190]

1059. Samsung's Accused Instrumentalities contain logic to store a value representing a size of the NAND flash memory that is to be physically erased in fulfilling a maintenance request issued by the host. '801 Patent, claim 1, element [1h].

1060. Samsung's Accused Instrumentalities are able to store a zone capacity, which is a value representing a size of the NAND flash memory. See:

[190] *Id*. at 10.

15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).

ZNS Specification, Section 3.4.2.2.3[191]

1061. Samsung's Accused Instrumentalities are able to physically erase the NAND flash memory in a zone in fulfilling a zone management send command specifying a reset zone action (maintenance request) issued by the host. See:

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

ZNS Specification, Section 3.4.3[192]

[191] *Id*. at 28.

[192] *Id*. at 29.

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[193]

1062. Samsung's Accused Instrumentalities contain logic to provide to the host, responsive to a query received from the host, the information, as updated, and the stored value. '801 Patent, claim 1, element [1i].

1063. Samsung's Accused Instrumentalities are able to provide to the host, responsive to a zone management receive command specifying a report zones action, the write pointer and the zone capacity. See:

[193] *Id*. at 30–31.

3.4.2.1 Zone Receive Actions

The Zone Management Receive command Zone Receive Action field specifies what action to perform.

3.4.2.1.1 Report Zones

The Report Zones action returns the Report Zones data structure (refer to Figure 35). The Zone Descriptors of the Report Zones data structure shall:

a) report only Zone Descriptors of zones for which the ZSLBA value is greater than or equal to the ZSLBA value of the zone specified by the SLBA value in the command;
b) match the criteria in the Zone Receive Action Specific field; and
c) be sorted in ascending order by the ZSLBA value of each zone.

ZNS Specification, Section 3.4.2.1[194]

3.4.2.2.1 Report Zones Data Structure

Figure 35 defines the Report Zones Data Structure.

Figure 35: Report Zones Data Structure

Bytes	Description
07:00	**Number of Zones:** If the Partial Report bit (refer to Figure 34) is cleared to '0', then this field indicates the number of zones that match the criteria defined in section 3.4.2.1.1. If the Partial Report bit is set to '1', then this field indicates the number of zones for which complete Zone Descriptors were transferred to the data buffer. Refer to section 3.4.2.1.1 for the content of the data buffer.
63:08	Reserved
127:64	**Zone Descriptor 0:** Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).
191:128	**Zone Descriptor 1:** Contains the Zone Descriptor for the second zone reported, if any.
...	...
((n+1)*64)+63:(n+1)*64	**Zone Descriptor n:** Contains the Zone Descriptor for the last zone reported, if any.

ZNS Specification, Section 3.4.2.2.1[195]

[194] *Id*. at 26.

[195] *Id*.

3.4.2.2.3 Zone Descriptor Data Structure

Figure 37 defines the Zone Descriptor data structure.

Figure 37: Zone Descriptor Data Structure

Bytes	Description	
00	Bits	Description
	7:4	Reserved
	3:0	**Zone Type (ZT):** This field indicates the type of the zone.

Value	Definition
2h	Sequential Write Required
All other values	Reserved

15:08	**Zone Capacity (ZCAP):** This field contains the maximum number of logical blocks that are available to be written with user data when the zone is in the ZSE:Empty state. This value shall be less than or equal to the Zone Size field (refer to Figure 49). If the Variable Zone Capacity bit is cleared to '0' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure (refer to section 4.1.5.1), then this field does not change without a change to the format of the zoned namespace. If the Variable Zone Capacity bit is set to '1' in the Zone Operation Characteristics field in the Zoned Namespace Command Set specific Identify Namespace data structure, then the zone capacity may change upon successful completion of a Zone Management Send command specifying the Zone Send Action of Reset Zone (refer to section 3.4.3.1.4).
23:16	**Zone Start Logical Block Address (ZSLBA):** This field contains the 64-bit address of the lowest logical block for the zone.
31:24	**Write Pointer (WP):** This field is the logical block address where the next write operation for this zone should be issued. Refer to section 2.1.1.2.1 for the behavior of the write pointer.

ZNS Specification, Section 3.4.2.2.3[196]

1064. Samsung's Accused Instrumentalities contain logic to update metadata for the specific one of the respective subdivisions, wherein the metadata represents a characteristic of the data in the specific one of the respective subdivisions. '801 Patent, claim 1, element [1j].

1065. Samsung's Accused Instrumentalities update metadata for the zones, including metadata that represents a characteristic of the data in the zones. See:

[196] *Id*. at 27–28.



Samsung Over-Provisioning White Paper[197]

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-

[197] Ex. V, Samsung Over-Provisioning White Paper, at 1.

194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[198]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

[198] *Id*. at 6.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

Samsung SSD White Paper[199]

2. Why is OP important?

OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

[199] Ex. I, Samsung SSD White Paper, at 19.

Samsung Application Note[200]

5.4 Reset Zone Recommended

A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[201]

1066. Samsung's Accused Instrumentalities contain logic to, with respect to the metadata, as updated, perform a comparison based on the metadata to detect a condition. '801 Patent, claim 1, elements [1k], [1l].

1067. Samsung's Accused Instrumentalities perform a comparison based on the metadata as updated to detect a condition for maintenance purposes, including but not limited to a condition that certain threshold is met. See:



Samsung Over-Provisioning White Paper[202]

[200] Ex. CC, Samsung Application Note, at 2.

[201] Ex. Z, ZNS Specification, at 41.

[202] Ex. V, Samsung Over-Provisioning White Paper, at 1.

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-
194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[203]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

[203] *Id.* at 6.

Samsung SSD White Paper[204]

> **2. Why is OP important?**
>
> OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

Samsung Application Note[205]

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[206]

1068. Samsung's Accused Instrumentalities contain logic to, with respect to the metadata, as updated, store information, responsive to detection of the condition, wherein the information is accessible to the host and indicates that the condition was detected. '801 Patent, claim 1, elements [1k], [1m].

1069. ZNS Specification defines a "Reset Zone Recommended" field, which is information accessible to the host which indicates that certain condition was detected for a zone such that the memory device recommends the reset of that zone. See:

---

[204] Ex. I, Samsung SSD White Paper, at 19.

[205] Ex. CC, Samsung Application Note, at 2.

[206] Ex. Z, ZNS Specification, at 41.

5.4 Reset Zone Recommended

A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer to section 3.4.3.1.4). If a controller schedules such an internal operation on a zone, the controller may notify the host by:

a) setting the Reset Zone Recommended zone attribute of the specific zone to '1' (refer to Figure 37);
b) setting the Reset Zone Recommended Time Limit zone attribute information to indicate the number of seconds before the controller intends to perform an internal operation on the specified zone; and
c) generating a Zone Descriptor Changed event for the specific zone.

As a zone reset operation destroys data in a specific zone, it is optional for the host software to perform a zone reset operation on zones that have the Reset Zone Recommended zone attribute set to '1'. If the host does not perform a zone reset operation on the specific zone, then the internal operation, which may impact performance, may be performed.

If the controller has processed the internal operation or the internal operation is no longer scheduled, the controller may notify the host by:

a) clearing the Reset Zone Recommended zone attribute of the specific zone to '0'; and
b) generating a Zone Descriptor Changed event for the specific zone.

If a zone is in the ZSF:Read Only state or the ZSF:Offline state, then the Reset Zone Recommended attribute shall be cleared to '0'.

ZNS Specification, Section 5.4[207]

Figure 37: Zone Descriptor Data Structure

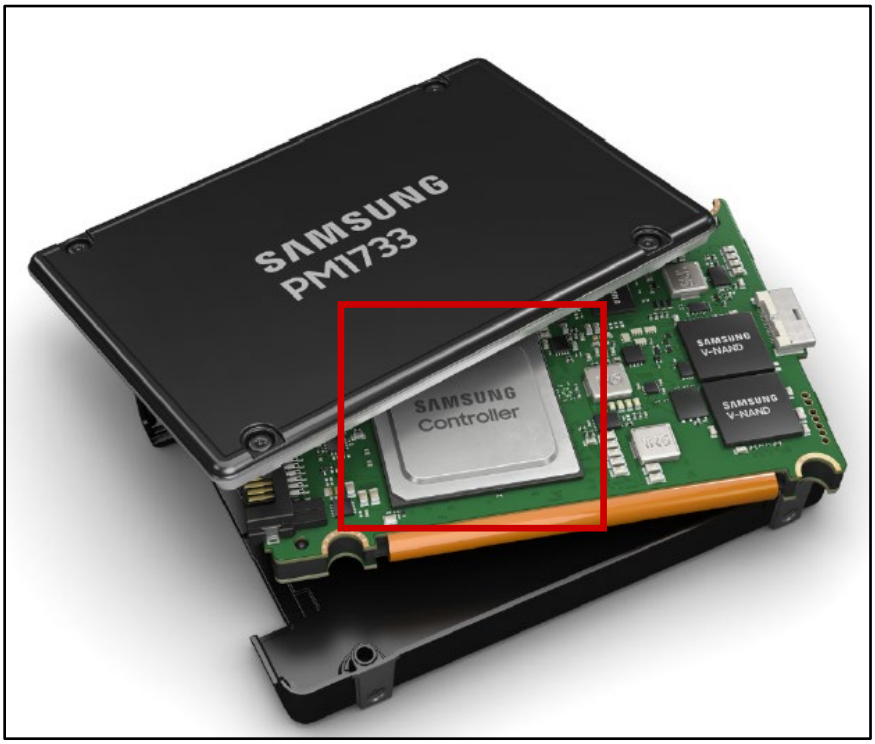
Bytes	Description
02	**Zone Attributes (ZA):** Indicates attributes for the Zone:

Bits	Description
7	**Zone Descriptor Extension Valid (ZDEV):** If this bit is set to '1', then Zone Descriptor Extension data is associated with the zone. If this bit is cleared to '0', then no Zone Descriptor Extension data is associated with the zone. Refer to section 5.1.
6:3	Reserved
2	**Reset Zone Recommended (RZR):** If this bit is set to '1', then the controller recommends that this zone be reset. Refer to section 5.4.
1	**Finish Zone Recommended (FZR):** If this bit is set to '1', then the controller recommends that this zone be finished. Refer to section 5.5.
0	**Zone Finished by Controller (ZFC):** If this bit is set to '1', then the controller finished this zone due to a Zone Active Excursion. Refer to section 5.6.

[207] *Id*. at 41–42.

ZNS Specification, Section 3.4.2.2.3[208]

1070. In Samsung's Accused Instrumentalities, each said logic comprises at least one of circuitry or instructions stored on physical storage media that, when executed, are to control operation of circuitry. '801 Patent, claim 1, element [1n]. See, for example:



Samsung PM1733 Product Brief[209]

1071. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '801 Patent.

1072. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

---

[208] *Id*. at 28.

[209] Ex. Q, Samsung PM1733 Product Brief (annotated).

1073. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1074. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1075. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

1076. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1077. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1078. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1079. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1080. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1081. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1082. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1083. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1084. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

1085. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1086. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1087. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1088. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1089. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1090. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

1091. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1092. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1093. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a

material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1094. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially

made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1095. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1096. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1097. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1098. SEA induces Samsung's customers' infringement by offering for sale, selling,

encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

1099. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

1100. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1101. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice

one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1102. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1103. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1104. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

1105. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

1106. SEC's infringement of the Asserted Patent is ongoing.

1107. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1108. SEC's ongoing infringement of the Asserted Patent is willful.

1109. SEC's ongoing infringement of the Asserted Patent is egregious.

1110. SEA's infringement of the Asserted Patent is ongoing.

1111. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1112. SEA's ongoing infringement of the Asserted Patent is willful.

1113. SEA's ongoing infringement of the Asserted Patent is egregious.

1114. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

1115. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

1116. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

1117. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

1118. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

1119. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT FIVE: INFRINGEMENT OF THE '657 PATENT

1120. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

1121. U.S. Patent No. 11,347,657 ("the '657 Patent"), entitled "Addressing Techniques for Write and Erase Operations in a Non-Volatile Storage Device," was legally and duly issued on May 31, 2022, naming Robert Lercari, Alan Chen, Mike Jadon, Craig Robertson, and Andrey V. Kuzmin and as the inventor. *See* Exhibit E.

1122. Radian owns all rights, title, and interest in the '657 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future infringement.

1123. Radian has complied with 35 U.S.C. § 287 with respect to the '657 Patent.

1124. The '657 Patent is related to improvements in memory addressing to provide flexibility and less latency. See, for example:

> This disclosure provides techniques hierarchical address virtualization within a memory controller and configurable block device allocation. By performing address translation only at select hierarchical levels, a memory controller can be designed to have predictable I/O latency, with brief or otherwise negligible logical-to-physical address translation time.[210]

1125. The '657 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> A memory controller that subdivides an incoming memory address into multiple discrete address fields corresponding to respective hierarchical groups of structural elements within a target nonvolatile semiconductor memory system and in which at least one of the discrete address fields constitutes a virtual address for the corresponding physical element within the structural hierarchy is disclosed in various embodiments. Through this hierarchical subdivision, the virtual address portion of the incoming memory

[210] Ex. E, U.S. Patent No. 11,347,657, at Abstract.

> address is ensured to resolve to an element within the physical bounds of a larger (hierarchically-superior) structure, but may be freely mapped to any of the constituent physical elements of that larger structure. Accordingly, a host requestor may issue logical memory addresses with address fields purposefully specified to direct read, write and maintenance operations to physically distinct structures within the memory system in a manner that limits performance-degrading conflicts while the memory controller remains free, by virtue of one or more virtualized address fields within the incoming logical addresses, to virtualize localized groups of physical structures and thus mask defective structural elements and swap operational structural elements into and out of service, for example, as they wear or otherwise require maintenance.[211]

1126. Some of these described solutions are reflected in the claimed inventions of the ’657 Patent.

1127. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the ’657 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the ’657 Patent.

1128. Claim 1 of the ’657 Patent recites:

> [preamble] A storage device comprising:
>
> [1a] flash memory comprising erase units; and
>
> [1b] circuitry to receive from a host, and control execution of, incoming requests including data write requests and erase requests, wherein each of the incoming requests is accompanied by incoming address information;
>
> [1c] wherein said circuitry is to, for each of the incoming requests:
>
> [1d] derive, from the accompanying incoming address information, a first address portion and a second address portion

[211] *Id*. at 2:66–3:21.

[1e] identify an addressed block device from the first address portion,

[1f] use a first division operation on the second address portion, to identify an addressed segment within the addressed block device, wherein the addressed segment corresponds to a subset of the erase units of the flash memory, and

[1g] for each of the data write requests, identify from the first division operation a third address portion, and use a second division operation on the third address portion to identify an address value to select one of the erase units from the subset which corresponds to the addressed segment and to identify a storage location within the selected one of the erase units;

[1h] wherein said circuitry is to execute each one of the data write requests by programming accompanying data into the storage location that was identified for that one of the data write requests; and

[1i] wherein said circuitry is also to execute each one of the incoming erase requests by controlling erasure in the flash memory of one or more of the erase units in the subset corresponding to the addressed segment that was identified from the incoming address information accompanying the one of the incoming erase requests.

1129. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

1130. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1131. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1132. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

1133. Changes from one revision to another in the cited NVMe specifications are not

material to infringement.

1134. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

1135. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a storage device. '657 Patent, claim 1, preamble.

1136. Samsung's Accused Instrumentalities contain a flash memory comprising erase units. '657 Patent, claim 1, element [1a].

1137. For example, Samsung's PM1733 contains NAND flash memory. See:

Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[212]

1138. The flash memory in Samsung's Accused Instrumentalities comprises blocks (erase units). See:

[212] Ex. Q, Samsung PM1733 Product Brief (annotated).

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[213]

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Samsung Over-Provisioning White Paper[214]

– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[215]

[213] Ex. I, Samsung SSD White Paper, at 17.

[214] Ex. V, Samsung Over-Provisioning White Paper, at 2 (highlighted).

[215] Ex. W, FLASH MEMORY INTEGRATION, at 22 (highlighted).

> **2.2.1. *Erase-before-write constraint***
>
> Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[216]

1139. Samsung's Accused Instrumentalities contain circuitry to receive from a host, and control execution of, incoming requests including data write requests and erase requests, wherein each of the incoming requests is accompanied by incoming address information. '657 Patent, claim 1, element [1b].

1140. Samsung's Accused Instrumentalities are able to receive and execute a Write command from a host.

1141. A Write command is a data write request.

1142. A Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. See:

> **3.3.6 Write command**
>
> The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

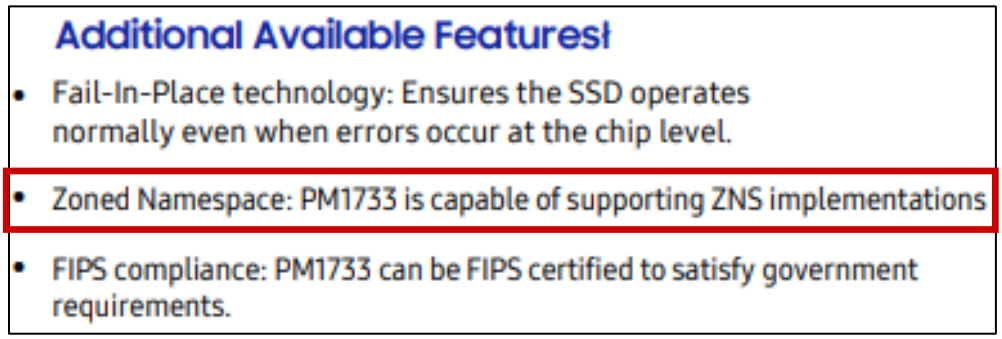
NVM Command Set Specification, Section 3.3.6[217]

---

[216] *Id*. at 24.

[217] Ex, AA, NVM Command Set Specification, at 53.

1143. The Write command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:



Samsung PM1733 Product Brief[218]

[218] Ex. Q, Samsung PM1733 Product Brief (annotated).

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[219]

[219] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[220]

1144. The NSID is incoming address information.

1145. The Write command is accompanied by a logical block address ("LBA"). See:

[220] Ex. AA, NVMe Base Specification, at 107.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2

ZNS Specification, Section 3.2[221]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[222]

1146. The starting LBA is incoming address information.

1147. Samsung's Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action from the host. See:

[221] Ex. Z, ZNS Specification, at 18 (annotated).

[222] Ex. AA, NVM Command Set Specification, at 54.

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

ZNS Specification, Section 3.4.3[223]

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

[223] Ex. Z, ZNS Specification, at 29.

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[224]

1148. A Zone Management Send command specifying a Reset Zone action is an erase request.

1149. The Zone Management Send command is accompanied by the starting LBA of a zone. See:

3.4.3 Zone Management Send command

The Zone Management Send command requests an action on one or more zones. The command uses the Data Pointer, Command Dword 10, Command Dword 11, and Command Dword 13 fields. All other command specific fields are reserved.

Figure 38: Zone Management Send – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field specifies the lowest LBA of the zone on which the Zone Send Action is performed. Command Dword 10 contains bits 31:00 of the SLBA; Command Dword 11 contains bits 63:32 of the SLBA.

[224] *Id*. at 30–31.

ZNS Specification, Section 3.4.3[225]

1150. The starting LBA is incoming address information.

1151. Samsung's Accused Instrumentalities satisfy the claim element "wherein said circuitry is to, for each of the incoming requests." '657 Patent, claim 1, element [1c]. Samsung's Accused Instrumentalities contain said circuitry that is to, for each of the incoming requests, perform elements [1d] through [1g]. See below.

1152. Samsung's Accused Instrumentalities contain said circuitry to derive, from the accompanying incoming address information, a first address portion and a second address portion. '657 Patent, claim 1, element [1d].

1153. Samsung's Accused Instrumentalities are able to derive, from the accompanying incoming address information, a Namespace Identifier or NSID (a first address portion) and a starting LBA (a second address portion). See:

[225] *Id*. at 29.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[226]

[226] *Id*. at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[227]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[228]

Figure 38: Zone Management Send – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field specifies the lowest LBA of the zone on which the Zone Send Action is performed. Command Dword 10 contains bits 31:00 of the SLBA; Command Dword 11 contains bits 63:32 of the SLBA.

ZNS Specification, Section 3.4.3[229]

[227] Ex. BB, NVMe Base Specification, at 107.

[228] Ex. AA, NVM Command Set Specification, at 54.

[229] Ex. Z, ZNS Specification, at 29.

1154. Samsung's Accused Instrumentalities contain said circuitry to identify an addressed block device from the first address portion. '657 Patent, claim 1, element [1e].

1155. Samsung's Accused Instrumentalities are able to identify an addressed namespace (block device) from the Namespace Identifier or NSID (the first address portion). See:

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[230]

[230] *Id*. at 18.

Figure 88: Common Command Format

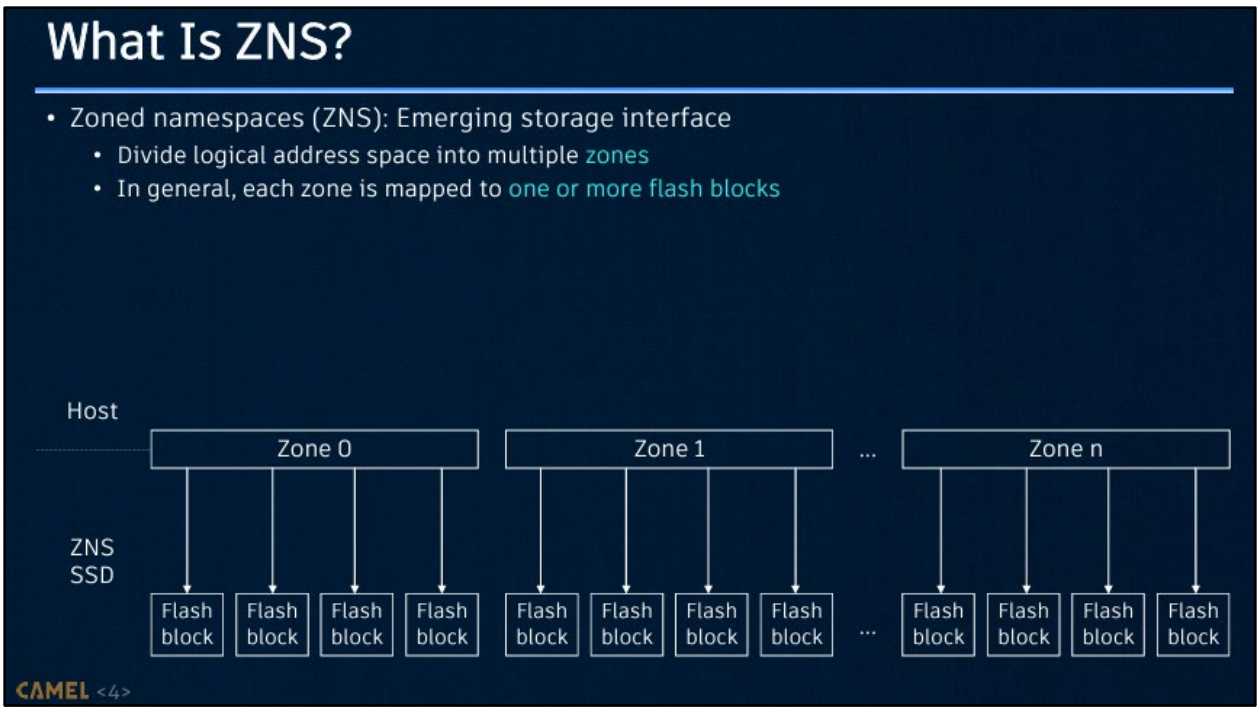
Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[231]

1156. Samsung's Accused Instrumentalities contain said circuitry to use a first division operation on the second address portion, to identify an addressed segment within the addressed block device, wherein the addressed segment corresponds to a subset of the erase units of the flash memory. '657 Patent, claim 1, element [1f].

1157. Samsung's Accused Instrumentalities contain zones (segments) that correspond to subsets of the erase units of the flash memory. See:

[231] Ex. BB, NVMe Base Specification, at 107.



Presentation on ZNS[232]

> This new attribute was introduced to allow for the zone size to remain a power of two number of logical blocks (facilitating logical block to zone number conversions) while allowing optimized mapping of a zone storage capacity to the underlying media characteristics. For instance, in the case a flash based device, a zone capacity can be aligned to the size of flash erase blocks without requiring that the device implements a power-of-two sized erased block.

SSDs with NVMe Zoned Namespace (ZNS) Support[233]

[232] Ex. X, *What You Can't Forget: Exploiting Parallelism for Zoned Namespaces*, at 4.

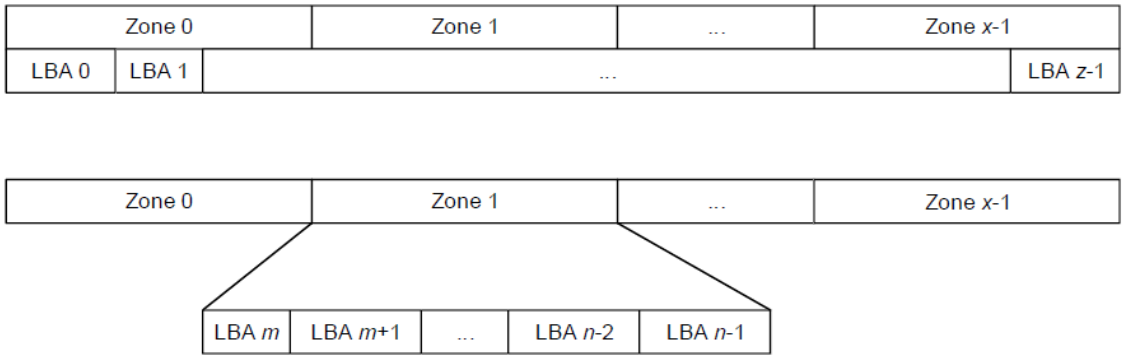
[233] Ex. Y, *SSDs with NVMe Zoned Namespace (ZNS) Support*.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[234]

1158. Samsung's Accused Instrumentalities are able to use a first division operation on the starting LBA to identify an addressed zone (segment) within the addressed namespace (block device). All zones are equally sized with contiguous non-overlapping ranges of LBAs. See:

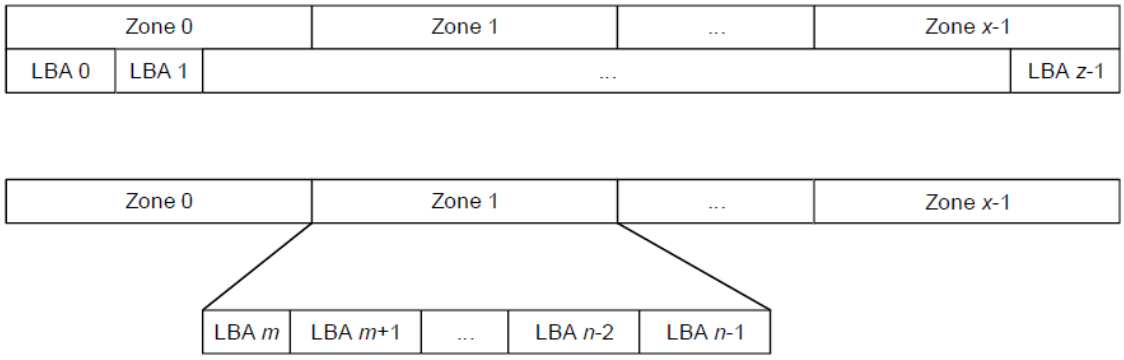
[234] Ex. Z, ZNS Specification, at 8.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[235]

[235] *Id*. at 8.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[236]

[236] *Id*. at 20.

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

If the command completes successfully, depending on the Zone Send Action field and the current states of the zones specified by the command, then that command may affect zones in various ways, including the following:

a) the zone state may change;
b) the Zone Descriptor Extension data may change; and
c) the Zone Descriptor Extension Valid zone attribute bit may change.

ZNS Specification, Section 3.4.3[237]

Figure 40: Zone Management Send – Command Specific Status Values

Value	Description
BAh	**Zone Is Read Only:** Zone is in the ZSRO:Read Only state. This may have occurred during the processing of the command.
BBh	**Zone Is Offline:** Zone is in the ZSO:Offline state. This may have occurred during the processing of the command.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.
BFh	**Invalid Zone State Transition:** The request is not a valid zone state transition.

ZNS Specification, Section 3.4.3.2[238]

[237] *Id*. at 29.

[238] *Id*. at 32.

1159. Samsung's Accused Instrumentalities contain said circuitry to, for each of the data write requests, identify from the first division operation a third address portion, and use a second division operation on the third address portion to identify an address value to select one of the erase units from the subset which corresponds to the addressed segment and to identify a storage location within the selected one of the erase units. '657 Patent, claim 1, element [1g].

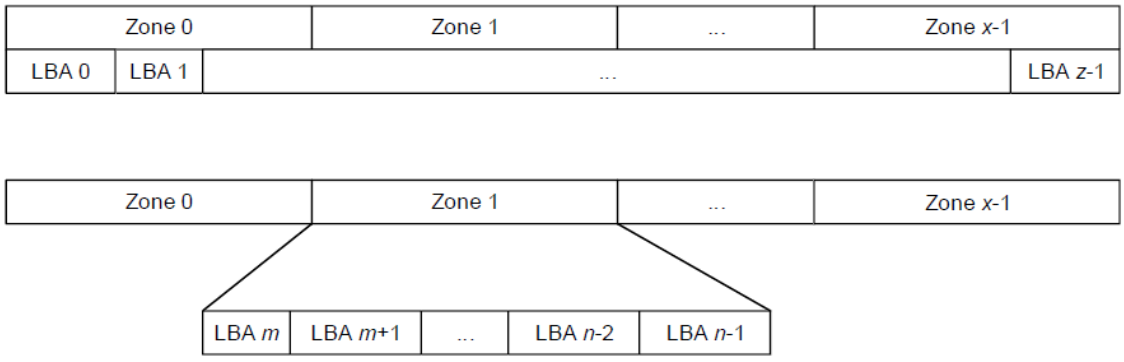
1160. Samsung's Accused Instrumentalities are able to identify from the first division operation a third address portion, which is the calculated remainder of said division operation. Samsung's Accused Instrumentalities are able to use such remainder to identify an address value to select one of the erase units corresponding to the addressed zone, and to identify a storage location within such erase unit. See:

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[239]

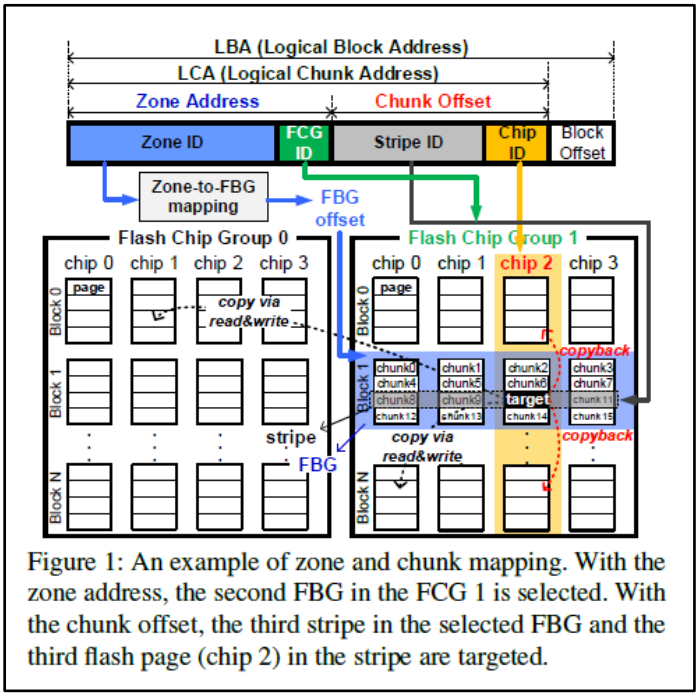


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[240]

1161. Samsung's Accused Instrumentalities contain said circuitry to execute each one of the data write requests by programming accompanying data into the storage location that was identified for that one of the data write requests. '657 Patent, claim 1, element [1h].

1162. Samsung's Accused Instrumentalities are able to execute data write requests by programming data into the identified storage location. See:

[239] *Id*. at 8.

[240] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[241]

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

[241] Ex. Z, ZNS Specification, at 20.

NVMe Base Specification, Section 3.3.3.1[242]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[243]

1163. Samsung's Accused Instrumentalities contain said circuitry to execute each one of the incoming erase requests by controlling erasure in the flash memory of one or more of the erase units in the subset corresponding to the addressed segment that was identified from the incoming address information accompanying the one of the incoming erase requests. '657 Patent, claim 1, element [1i].

1164. Samsung's Accused Instrumentalities are able to execute the Zone Management Send command specifying a Reset Zone action by controlling erasure of one or more of the erase units corresponding to the addressed zone. See:

[242] Ex. BB, NVMe Base Specification, at 107.

[243] Ex. AA, NVM Command Set Specification, at 54.

Figure 39: Zone Management Send – Command Dword 13

Bits	Description
31:09	Reserved
08	**Select All:** If this bit is set to '1', then the SLBA field is ignored. If this bit is cleared to '0', then the SLBA field specifies the lowest logical block of the zone. Refer to section 3.4.3.1 for specific behavior for each Zone Send Action.
07:00	**Zone Send Action (ZSA):** Defines the zone action to be performed for Zone Management Send.

Value	Description	Refer to section
00h	Reserved	
01h	**Close Zone:** Close one or more zones.	3.4.3.1.1
02h	**Finish Zone:** Finish one or more zones.	3.4.3.1.2
03h	**Open Zone:** Open one or more zones.	3.4.3.1.3
04h	**Reset Zone:** Reset one or more zones.	3.4.3.1.4
05h	**Offline Zone:** Offline one or more zones.	3.4.3.1.5
06h to 0Fh	Reserved	
10h	**Set Zone Descriptor Extension:** Attach Zone Descriptor Extension data to a zone.	3.4.3.1.6
11h to FFh	Reserved	

ZNS Specification, Section 3.4.3[244]

3.4.3.1.4 Reset Zone

If the Select All bit in Command Dword 13 is cleared to '0', and the zone specified by the SLBA field is in the:

a) ZSIO:Implicitly Opened state, the ZSEO:Explicitly Opened state, the ZSC:Closed state, or the ZSF:Full state, then the specified zone shall be transitioned to the ZSE:Empty state;

[244] Ex. Z, ZNS Specification, at 29.

b) ZSE:Empty state, then no change shall be made to the zone state; and
c) ZSRO:Read Only state, or the ZSO:Offline state, then the controller shall abort the command with a status code of Invalid Zone State Transition.

If the Select All bit is set to '1', then the SLBA field shall be ignored and all zones that are in the:

a) ZSIO:Implicity Opened state;
b) ZSEO:Explicitly Opened state;
c) ZSC:Closed state; and
d) the ZSF:Full state,

shall be transitioned to the ZSE:Empty state.

If the command completes successfully, then for each affected zone:

a) the Write Pointer zone attribute in the Zone Descriptor shall be set to the ZSLBA of the zone; and
b) the following zone attribute bits in the Zone Descriptor shall be cleared to '0':

a) Zone Descriptor Extension Valid;
b) Finish Zone Recommended;
c) Reset Zone Recommended; and
d) Zone Finished by Controller.

ZNS Specification, Section 3.4.3.1.4[245]

1165. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '657 Patent.

1166. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1167. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1168. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1169. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

1170. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities

[245] *Id*. at 30–31.

in the United States.

1171. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1172. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1173. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1174. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1175. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1176. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1177. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1178. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

1179. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1180. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1181. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1182. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1183. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1184. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

1185. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1186. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1187. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their

components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1188. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1189. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice

one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1190. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1191. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1192. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

1193. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

1194. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1195. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the

Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1196. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1197. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1198. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

1199. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

1200. SEC's infringement of the Asserted Patent is ongoing.

1201. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1202. SEC's ongoing infringement of the Asserted Patent is willful.

1203. SEC's ongoing infringement of the Asserted Patent is egregious.

1204. SEA's infringement of the Asserted Patent is ongoing.

1205. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1206. SEA's ongoing infringement of the Asserted Patent is willful.

1207. SEA's ongoing infringement of the Asserted Patent is egregious.

1208. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

1209. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

1210. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

1211. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

1212. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

1213. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT SIX: INFRINGEMENT OF THE '656 PATENT

1214. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

1215. U.S. Patent No. 11,347,656 ("the '656 Patent"), entitled "Storage Drive with Geometry Emulation Based on Division Addressing and Decoupled Bad Block Management," was legally and duly issued on May 31, 2022, naming Robert Lercari, Alan Chen, Mike Jadon, Craig Robertson, and Andrey V. Kuzmin and as the inventor. *See* Exhibit F.

1216. Radian owns all rights, title, and interest in the '656 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future

infringement.

1217. Radian has complied with 35 U.S.C. § 287 with respect to the ’656 Patent.

1218. The ’656 Patent is related to improvements in memory addressing to provide flexibility and less latency. See, for example:

> This disclosure provides techniques hierarchical address virtualization within a memory controller and configurable block device allocation. By performing address translation only at select hierarchical levels, a memory controller can be designed to have predictable I/O latency, with brief or otherwise negligible logical-to-physical address translation time.[246]

1219. The ’656 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> A memory controller that subdivides an incoming memory address into multiple discrete address fields corresponding to respective hierarchical groups of structural elements within a target nonvolatile semiconductor memory system and in which at least one of the discrete address fields constitutes a virtual address for the corresponding physical element within the structural hierarchy is disclosed in various embodiments. Through this hierarchical subdivision, the virtual address portion of the incoming memory address is ensured to resolve to an element within the physical bounds of a larger (hierarchically-superior) structure, but may be freely mapped to any of the constituent physical elements of that larger structure. Accordingly, a host requestor may issue logical memory addresses with address fields purposefully specified to direct read, write and maintenance operations to physically distinct structures within the memory system in a manner that limits performance-degrading conflicts while the memory controller remains free, by virtue of one or more virtualized address fields within the incoming logical addresses, to virtualize localized groups of physical structures and thus mask defective structural elements

[246] Ex. F, U.S. Patent No. 11,347,656, at Abstract.

> and swap operational structural elements into and out of service, for example, as they wear or otherwise require maintenance.[247]

1220. Some of these described solutions are reflected in the claimed inventions of the '656 Patent.

1221. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the '656 Patent by, among other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the '656 Patent.

1222. Claim 1 of the '656 Patent recites:

> [preamble] A storage device comprising:
>
> [1a] flash memory having physical erase units; and
>
> [1b] circuitry to receive incoming data access requests and for each of the incoming data access requests, accompanying address information, wherein said circuitry is to:
>
> [1c] derive, from the accompanying address information, a first address portion, a second address portion, and a third address portion,
>
> [1d] identify an addressed block device from the first address portion,
>
> [1e] identify an addressed segment within the addressed block device from the second address portion, and
>
> [1f] identify an addressed storage location within the addressed segment from the third address portion,
>
> [1g] wherein the circuitry is to use a division operation to identify a logical erase unit within the addressed segment,

---

[247] *Id*. at 2:66–3:21.

> [1h] wherein the circuitry is to identify one of the physical erase units that corresponds to the identified logical erase unit;
>
> [1i] wherein the storage device is to detect a failure condition of the physical erase units and is to remap each logical erase unit which corresponds to one of the physical erase units for which the failure condition is detected to instead correspond to one of the physical erase units for which the failure condition has not been detected and is to utilize the remapped correspondence to service one or more of the incoming data access requests.

1223. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

1224. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1225. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1226. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

1227. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

1228. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

1229. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a storage device. '656 Patent, claim 1, preamble.

1230. Samsung's Accused Instrumentalities contain a flash memory having physical erase units. '656 Patent, claim 1, element [1a].

1231. For example, Samsung's PM1733 contains NAND flash memory. See:

Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[248]

1232. The flash memory in Samsung's Accused Instrumentalities has physical blocks (erase units). See:

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[249]

[248] Ex. Q, Samsung PM1733 Product Brief (annotated).

[249] Ex. I, Samsung SSD White Paper, at 17.

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

Samsung Over-Provisioning White Paper[250]

– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[251]

2.2.1. *Erase-before-write constraint*

Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[252]

1233. Samsung's Accused Instrumentalities contain circuitry to receive incoming data access requests and for each of the incoming data access requests, accompanying address information. '656 Patent, claim 1, element [1b].

---

[250] Ex. V, Samsung Over-Provisioning White Paper, at 2 (highlighted).

[251] Ex. W, FLASH MEMORY INTEGRATION, at 22 (highlighted).

[252] *Id*. at 24.

1234. Samsung's Accused Instrumentalities are able to receive incoming Read commands from a host.

1235. A Read command is a data access request.

1236. A Read command reads data and metadata, if applicable, from the I/O controller for the LBAs indicated. See:

> **3.3.4 Read command**
>
> The Read command reads data and metadata, if applicable, from the I/O controller for the LBAs indicated. The command may specify protection information to be checked as part of the read operation.

NVM Command Set Specification, Section 3.3.4[253]

1237. The Read command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:

[253] Ex. AA, NVM Command Set Specification, at 48.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[254]

[254] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[255]

1238. The NSID is incoming address information.

1239. The Read command is accompanied by a logical block address ("LBA"). See:

[255] Ex. BB, NVMe Base Specification, at 107.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[256]

Figure 52: Read – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be read as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63: 32.

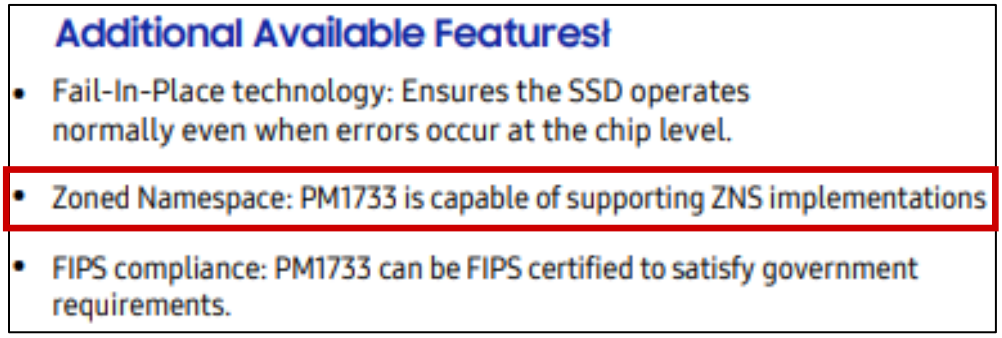
[256] Ex. Z, ZNS Specification, at 18.

NVM Command Set Specification, Section 3.3.4[257]

1240. The starting LBA is address information.

1241. Samsung's Accused Instrumentalities contain said circuitry to derive, from the accompanying address information, a first address portion, a second address portion and a third address portion. '656 Patent, claim 1, element [1c].

1242. Samsung's Accused Instrumentalities are able to derive, from the accompanying address information, a Namespace Identifier or NSID (a first address portion), a zone identifier (a second address portion), and an LBA and/or alternatively a zone offset (a third address portion). See:



Samsung PM1733 Product Brief[258]

[257] Ex. AA, NVM Command Set Specification, at 49.

[258] Ex. Q, PM1733 Product Brief (annotated).

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[259]

[259] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[260]

Figure 52: Read – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be read as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63: 32.

NVM Command Set Specification, Section 3.3.4[261]

[260] Ex. BB, NVMe Base Specification, at 107.

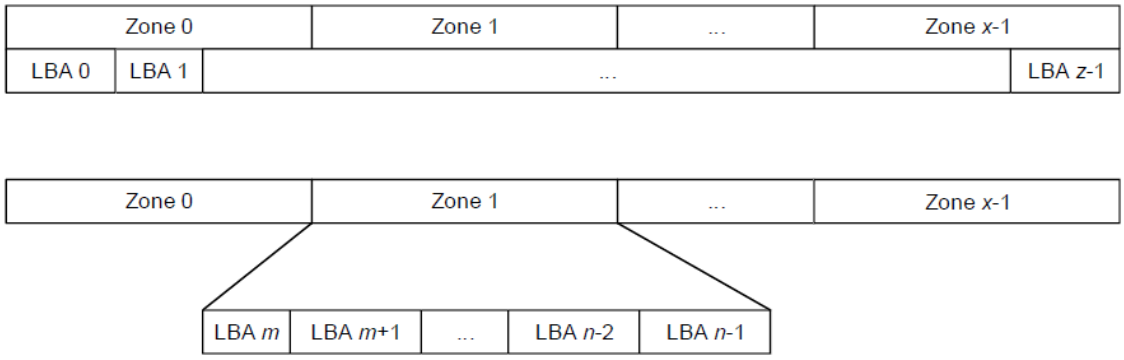
[261] Ex. AA, NVM Command Set Specification, at 49.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[262]

3.3.5 Read command

The Read command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.5.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 17.

Figure 17: Read – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.

[262] Ex. Z, ZNS Specification, at 8.

ZNS Specification, Section 3.3.5[263]

1243. Samsung's Accused Instrumentalities contain said circuitry to identify an addressed block device from the first address portion. '656 Patent, claim 1, element [1d].

1244. Samsung's Accused Instrumentalities are able to identify an addressed namespace (block device) from the Namespace Identifier or NSID (the first address portion). See:

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[264]

1245. Samsung's Accused Instrumentalities contain said circuitry to identify an addressed segment within the addressed block device from the second address portion. '656 Patent, claim 1, element [1e].

1246. Samsung's Accused Instrumentalities are able to identify an addressed zone

[263] *Id*. at 20.

[264] Ex. BB, NVMe Base Specification, at 107.

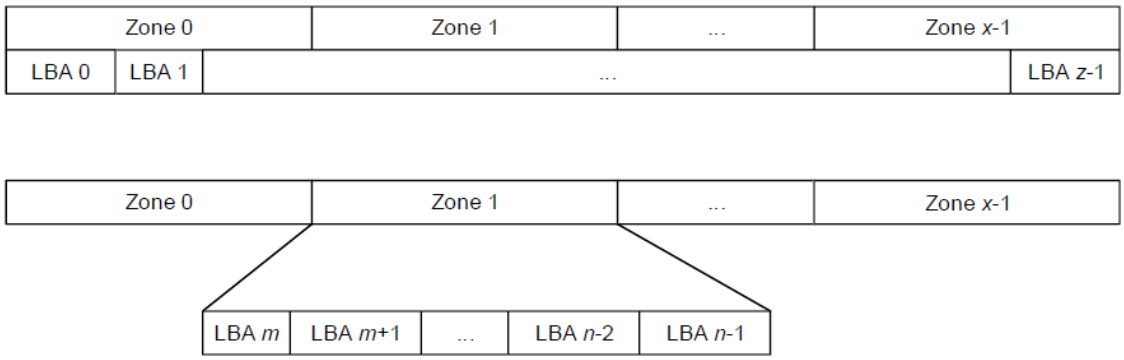
(segment) within the addressed namespace from the second address portion. See:

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[265]

[265] Ex. Z, ZNS Specification, at 8.

3.3.5 Read command

The Read command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.5.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 17.

Figure 17: Read – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.

ZNS Specification, Section 3.3.5[266]

1247. Samsung's Accused Instrumentalities contain said circuitry to identify an addressed storage location within the addressed segment from the third address portion. '656 Patent, claim 1, element [1f].

1248. Samsung's Accused Instrumentalities are able to identify an addressed storage location within the addressed zone from the LBA and/or alternatively a zone offset (the third address portion). See:

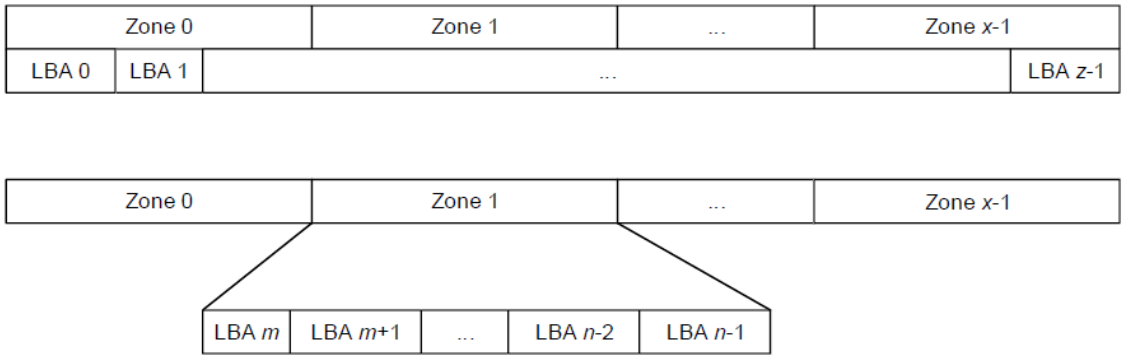
[266] *Id*. at 20.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[267]

1.4.3.3 logical block

The smallest addressable data unit for Read and Write commands.

1.4.3.4 logical block address (LBA)

The address of a logical block, referred to commonly as LBA.

NVM Command Set Specification, Section 1.4[268]

[267] *Id*. at 8.

[268] Ex. AA, NVM Command Set Specification, at 10.

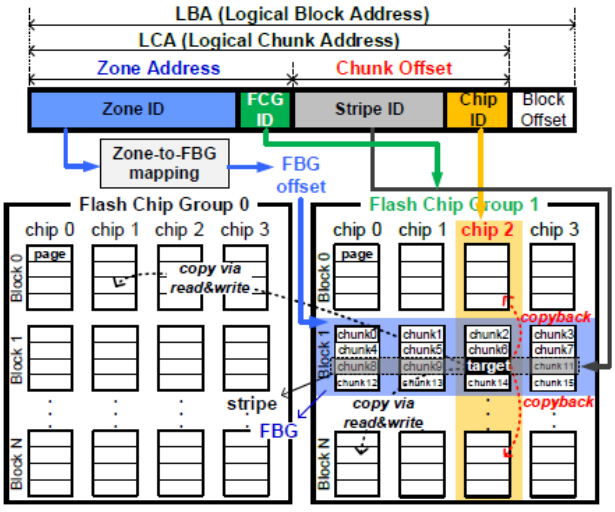


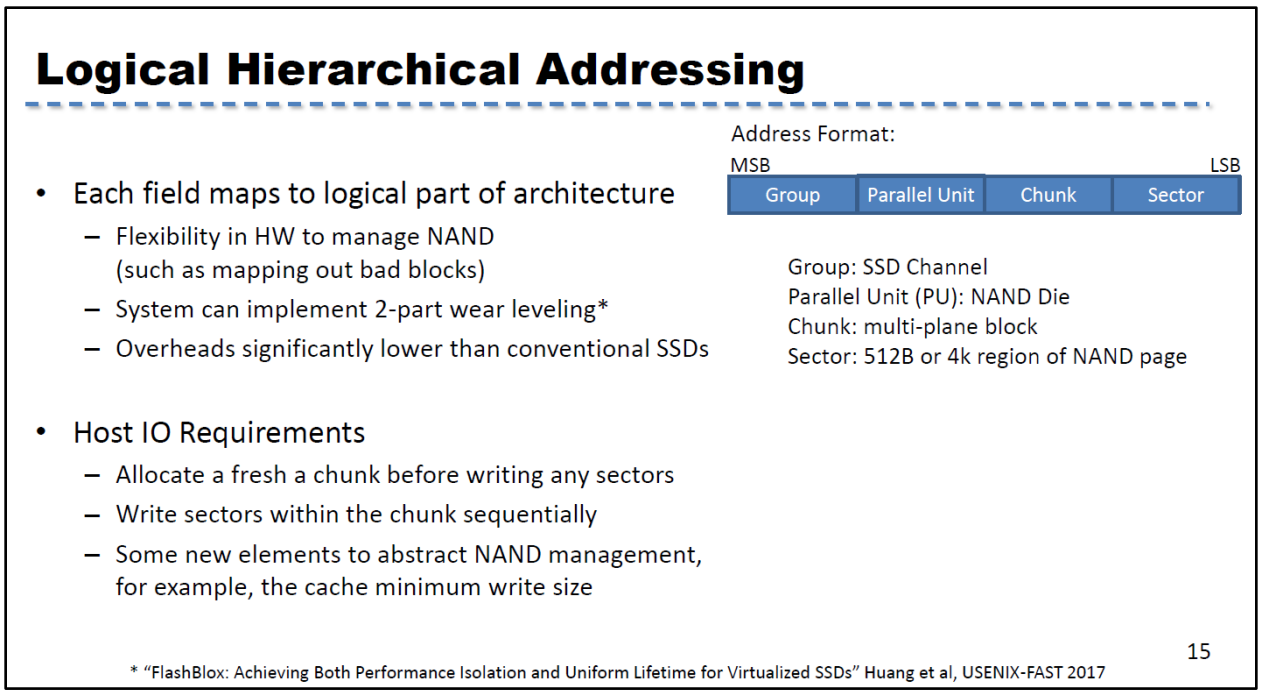
Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[269]

1249. Samsung's Accused Instrumentalities contain said circuitry to use a division operation to identify a logical erase unit within the addressed segment. '656 Patent, claim 1, element [1g].

1250. Samsung's Accused Instrumentalities are able to use a division operation to identify a logical erase unit within the addressed zone. See:

[269] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.



Microsoft Denali Presentation[270]

[270] Ex. DD, Microsoft Denali Presentation, at 15.

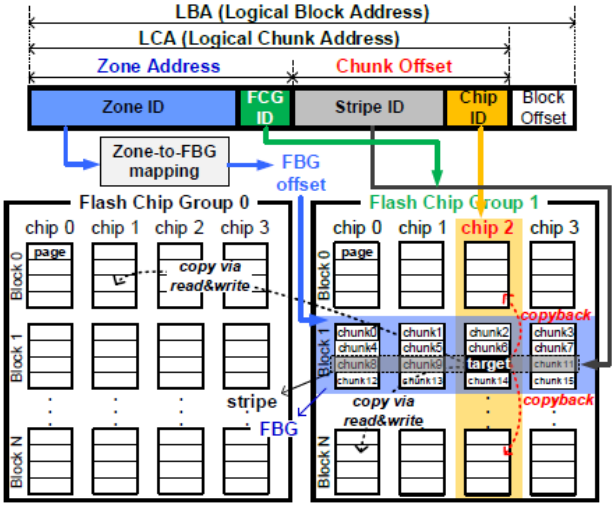
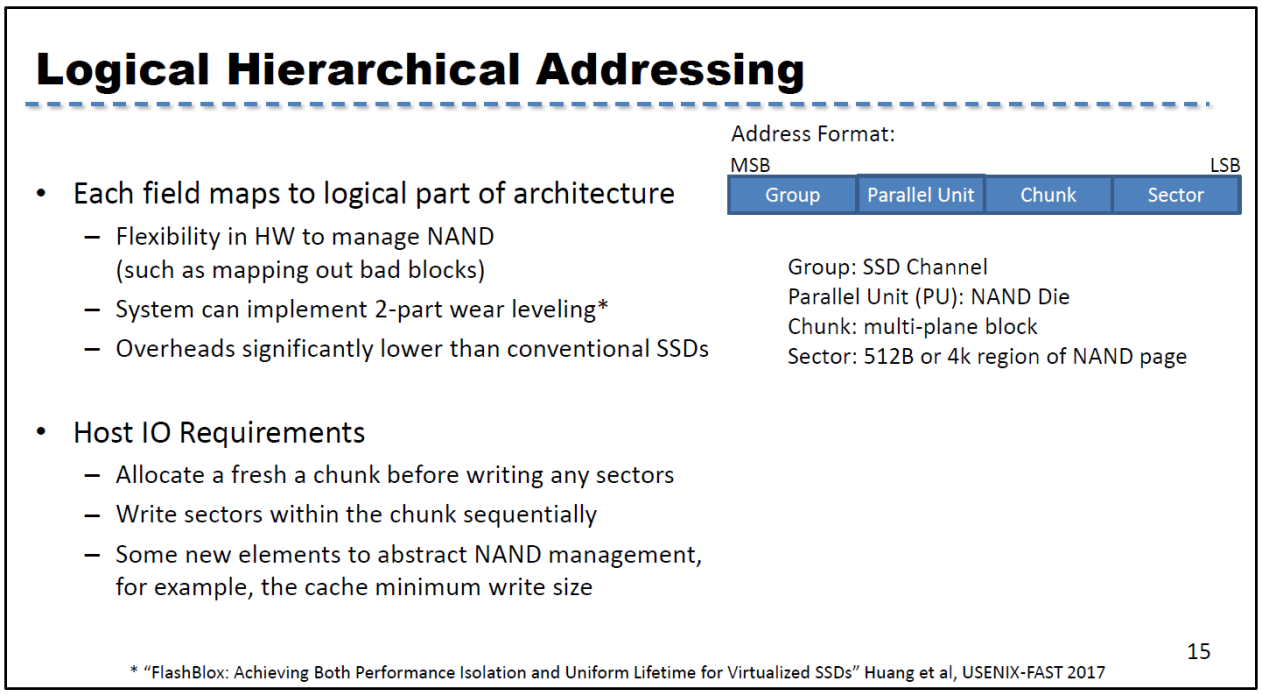


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[271]

1251. Samsung's Accused Instrumentalities contain said circuitry to identify one of the physical erase units that corresponds to the identified logical erase unit. '656 Patent, claim 1, element [1h].

[271] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.



Microsoft Denali Presentation[272]

[272] Ex. DD, Microsoft Denali Presentation, at 15.

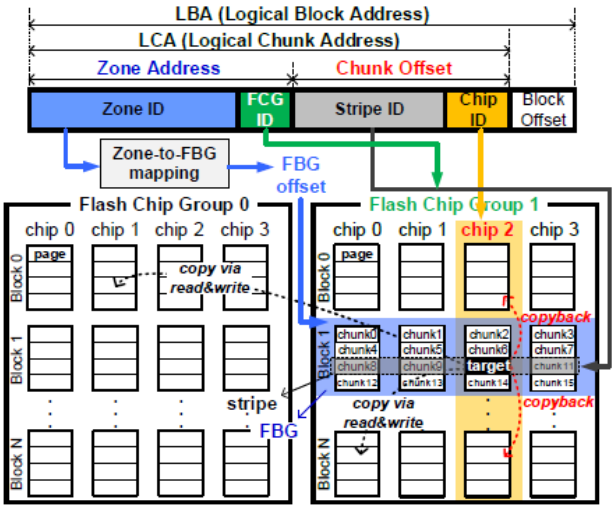


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[273]

1252. Samsung's Accused Instrumentalities detect a failure condition of the physical erase units and to remap each logical erase unit which corresponds to one of the physical erase units for which the failure condition is detected to instead correspond to one of the physical erase units for which the failure condition has not been detected and to utilize the remapped correspondence to service one or more of the incoming data access requests. '656 Patent, claim 1, element [1i].

1253. Samsung's Accused Instrumentalities detect a failure condition of the physical blocks, map out the bad blocks, and substitute different blocks for the bad ones. See:

[273] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

Samsung SSD White Paper[274]

2. Why is OP important?

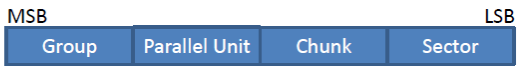
OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

Samsung Application Note[275]

Logical Hierarchical Addressing

- Each field maps to logical part of architecture
  - Flexibility in HW to manage NAND (such as mapping out bad blocks)
  - System can implement 2-part wear leveling*
  - Overheads significantly lower than conventional SSDs
- Host IO Requirements
  - Allocate a fresh a chunk before writing any sectors
  - Write sectors within the chunk sequentially
  - Some new elements to abstract NAND management, for example, the cache minimum write size

Address Format:



Group: SSD Channel
Parallel Unit (PU): NAND Die
Chunk: multi-plane block
Sector: 512B or 4k region of NAND page

[274] Ex. I, Samsung SSD White Paper, at 19.

[275] Ex. CC, Samsung Application Note, at 2.

Microsoft Denali Presentation[276]

1254. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '656 Patent.

1255. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1256. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1257. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1258. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

1259. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1260. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1261. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1262. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1263. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

---

[276] Ex. DD, Microsoft Denali Presentation, at 15.

1264. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1265. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1266. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1267. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

1268. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1269. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1270. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1271. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1272. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1273. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

1274. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building,

manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1275. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1276. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1277. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes

or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1278. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted

Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1279. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1280. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1281. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

1282. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

1283. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the

Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1284. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1285. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1286. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities'

infringement thereof, and the infringing nature of others' actions.

1287. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

1288. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

1289. SEC's infringement of the Asserted Patent is ongoing.

1290. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1291. SEC's ongoing infringement of the Asserted Patent is willful.

1292. SEC's ongoing infringement of the Asserted Patent is egregious.

1293. SEA's infringement of the Asserted Patent is ongoing.

1294. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1295. SEA's ongoing infringement of the Asserted Patent is willful.

1296. SEA's ongoing infringement of the Asserted Patent is egregious.

1297. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

1298. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

1299. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

1300. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

1301. SEC's actions have caused damage to Radian. Radian is entitled to recover from

SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

1302. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

COUNT SEVEN: INFRINGEMENT OF THE '995 PATENT

1303. Radian incorporates by reference the preceding paragraphs as if fully set forth herein.

1304. U.S. Patent No. 11,307,995 ("the '995 Patent"), entitled "Storage Device With Geometry Emulation Based on Division Programming and Decoupled NAND Maintenance," was legally and duly issued on April 19, 2022, naming Robert Lercari, Alan Chen, Mike Jadon, Craig Robertson, and Andrey V. Kuzmin and as the inventor. *See* Exhibit G.

1305. Radian owns all rights, title, and interest in the '995 Patent, and holds all substantial rights pertinent to this suit, including the right to sue and recover for all past, current, and future infringement.

1306. Radian has complied with 35 U.S.C. § 287 with respect to the '995 Patent.

1307. The '995 Patent is related to improvements in memory addressing to provide flexibility and less latency. See, for example:

> This disclosure provides techniques hierarchical address virtualization within a memory controller and configurable block device allocation. By performing address translation only at select hierarchical levels, a memory controller can be designed to have

> predictable I/O latency, with brief or otherwise negligible logical-to-physical address translation time.[277]

1308. The ’995 Patent proposed novel solutions that were not routine, ordinary, or conventional at the time of the inventions. Some exemplary embodiments of the inventions relating to these solutions are described in the specification. See, for example:

> A memory controller that subdivides an incoming memory address into multiple discrete address fields corresponding to respective hierarchical groups of structural elements within a target nonvolatile semiconductor memory system and in which at least one of the discrete address fields constitutes a virtual address for the corresponding physical element within the structural hierarchy is disclosed in various embodiments. Through this hierarchical subdivision, the virtual address portion of the incoming memory address is ensured to resolve to an element within the physical bounds of a larger (hierarchically-superior) structure, but may be freely mapped to any of the constituent physical elements of that larger structure. Accordingly, a host requestor may issue logical memory addresses with address fields purposefully specified to direct read, write and maintenance operations to physically distinct structures within the memory system in a manner that limits performance-degrading conflicts while the memory controller remains free, by virtue of one or more virtualized address fields within the incoming logical addresses, to virtualize localized groups of physical structures and thus mask defective structural elements and swap operational structural elements into and out of service, for example, as they wear or otherwise require maintenance.[278]

1309. Some of these described solutions are reflected in the claimed inventions of the ’995 Patent.

1310. On information and belief, Samsung directly infringed and is currently infringing, literally and/or under the doctrine of equivalents, at least one claim of the ’995 Patent by, among

---

[277] Ex. G, U.S. Patent No. 11,307,995, at Abstract.

[278] *Id*. at 2:66–3:21.

other things, making, using, selling, offering to sell, and/or importing within this District and elsewhere in the United States, without authority, the Accused Instrumentalities. For example, as shown below, the Accused Instrumentalities practice at least claim 1 of the ’995 Patent.

1311. Claim 1 of the ’995 Patent recites:

[preamble] A storage device comprising:

[1a] flash memory having physical erase units; and

[1b] circuitry to, for each of incoming data write requests:

[1c] derive, from incoming address information, a first address portion, a second address portion and a third address portion,

[1d] identify an addressed block device from the first address portion,

[1e] identify an addressed segment within the addressed block device from the second address portion, and

[1f] identify an addressed storage location within the addressed segment from the third address portion,

[1g] wherein the circuitry is to use a division operation on at least one of the incoming address information or information derived therefrom, to identify a logical erase unit within the addressed segment, and

[1h] wherein the circuitry is to identify one of the physical erase units that corresponds to the identified logical erase unit, and

[1i] program associated data into the identified one of the physical erase units;

[1j] wherein the storage device also comprises circuitry to:

[1k] maintain metadata for each physical erase unit which contains data, the metadata representing a time since data was programmed in the respective erase unit,

[1l] copy data from a first one of the physical erase units for which the maintained metadata meets at least one criterion to a second one of the physical erase units, and

[1m] remap a correspondence of a logical erase unit which

corresponds to the first one of the physical erase units to instead correspond to the second one of the physical erase units.

1312. Samsung's Accused Instrumentalities implement the NVMe Base Specification in the same ways or substantially the same ways that are material to infringement.

1313. Samsung's Accused Instrumentalities implement the NVMe Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1314. Samsung's Accused Instrumentalities implement the NVMe Zoned Namespace Command Set Specification in the same ways or substantially the same ways that are material to infringement.

1315. Samsung's Accused Instrumentalities implement other NVMe specifications in the same ways or substantially the same ways that are material to infringement.

1316. Changes from one revision to another in the cited NVMe specifications are not material to infringement.

1317. Samsung's Accused Instrumentalities infringe in the same way or substantially the same way. Variations in other features (e.g., form factor, housing, storage capacity, memory controller speeds, ports, or types of flash memory) are not material to infringement.

1318. To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities contain a storage device. '995 Patent, claim 1, preamble.

1319. Samsung's Accused Instrumentalities contain a flash memory having physical erase units. '995 Patent, claim 1, element [1a].

1320. For example, Samsung's PM1733 contains NAND flash memory. See:

Samsung PM1733 specifications

Form factor	U.2 / 2.5"
Capacity	1.92 TB, 3.84 TB, 7.68TB, and 15.36TB
Host interface	PCIe Gen 3/4 x4
Spec Compliance	NVMe spec rev. 1.3 PCI Express base specification rev. 4.0
NAND flash memory	Samsung V-NAND®

Samsung PM1733 Product Brief[279]

1321. The flash memory in Samsung's Accused Instrumentalities has physical blocks (erase units). See:

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a "page," which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a "block." While data is written in pages, it can only be erased in blocks. To help illustrate the concept

Samsung SSD White Paper[280]

Since overwriting is impossible with NAND flash, existing data must first be erased in order to write new data to that cell, which slows down overall write performance of the SSD. Generally, it takes longer to erase data than to write it because, as mentioned previously, write operations are carried out in pages while erase operations are executed in blocks. To alleviate this decrease in write performance, a process called garbage collection (GC) is implemented to create free blocks within the SSD. This technology secures free blocks by collecting valid pages into a single location and erasing the blocks consisting of invalid pages. However, this too may sometimes result in slower performance in the unexpected case that garbage collection interferes with the host write. Therefore, free space in the SSD is required to allow the firmware (FW) feature to run smoothly. This process in which extra space is allocated is called over-provisioning (OP).

[279] Ex. Q, Samsung PM1733 Product Brief (annotated).

[280] Ex. I, Samsung SSD White Paper, at 17.

Samsung Over-Provisioning White Paper[281]

– *Erase operation*, which is applied at the level of an entire block's granularity. Therefore, the erase affects all the pages contained in a given block.

Flash Memory Integration[282]

2.2.1. *Erase-before-write constraint*

Because of the internal architecture of a NAND chip, it is impossible to write data in a page that already contains data (*in-place* update). Before performing any write operation on a page which already contains data, this page has to be erased. However, the erase operation affects an entire block, and not just a single page. Moreover, this operation has a significant latency. In the literature, this limitation is sometimes referred to under the name of

Flash Memory Integration[283]

1322. Samsung's Accused Instrumentalities satisfy the claim element "circuitry to, for each of incoming data write requests." '995 Patent, claim 1, element [1b]. Samsung's Accused Instrumentalities contain circuitry to, for each of incoming data write requests, perform elements [1c] through [1i]. See below.

1323. Samsung's Accused Instrumentalities are able to receive and execute incoming Write commands from a host. See:

[281] Ex. V, Samsung Over-Provisioning White Paper, at 2 (highlighted).

[282] Ex. W, FLASH MEMORY INTEGRATION, at 22 (highlighted).

[283] *Id*. at 24.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[284]

1324. A Write command is a data write request.

1325. A Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. See:

[284] Ex. Z, ZNS Specification, at 18.

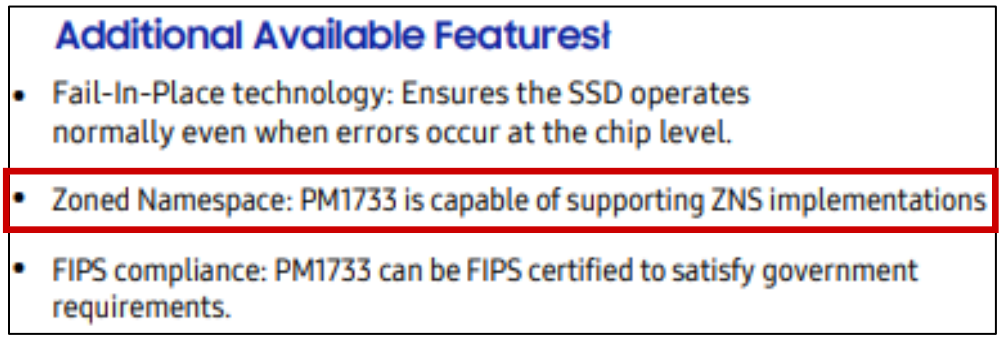
3.3.6 Write command

The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

NVM Command Set Specification, Section 3.3.6[285]

1326. Samsung's Accused Instrumentalities contain circuitry to derive, from incoming address information, a first address portion, a second address portion and a third address portion. '995 Patent, claim 1, element [1c].

1327. The Write command is accompanied by a Namespace Identifier ("NSID") that specifies the namespace to which the command applies. In the context of ZNS, the NSID specifies a zoned namespace to which the command applies. See:



Samsung PM1733 Product Brief[286]

[285] Ex. AA, NVM Command Set Specification, at 53.

[286] Ex. Q, Samsung PM1733 Product Brief (annotated).

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[287]

[287] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[288]

1328. The NSID is incoming address information.

1329. The Write command is accompanied by a logical block address ("LBA"). See:

[288] Ex. BB, NVMe Base Specification, at 107.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2

ZNS Specification, Section 3.2[289]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[290]

1330. The LBA is incoming address information.

1331. Samsung's Accused Instrumentalities are able to derive, from incoming address information, a Namespace Identifier or NSID (a first address portion), zone identifier information (a second address portion), and an LBA and/or alternatively a zone offset (a third address portion). See:

[289] Ex. Z, ZNS Specification, at 18 (annotated).

[290] Ex. AA, NVM Command Set Specification, at 54.

Figure 12: Opcodes for Zoned Namespace Command Set I/O Commands

Opcode by Field			Combined Opcode [1]	Command [2]	Reference
(07)	(06:02)	(01:00)			
Standard Command	Function	Data Transfer [3]			
Refer to the NVM Express Base Specification				Reservation Acquire	NVM Express Base Specification
Refer to the NVM Express Base Specification				Reservation Release	NVM Express Base Specification
NVM Command Set commands implemented by this specification					
Refer to the NVM Command Set Specification				Dataset Management	NVM Command Set Specification
NVM Command Set commands modified by this specification					
Refer to the NVM Command Set Specification				Write	3.3.7
Refer to the NVM Command Set Specification				Read	3.3.3
Refer to the NVM Command Set Specification				Write Uncorrectable	3.3.8
Refer to the NVM Command Set Specification				Compare	3.3.1
Refer to the NVM Command Set Specification				Write Zeroes	3.3.9
Refer to the NVM Command Set Specification				Verify	3.3.6
Refer to the NVM Command Set Specification				Copy	3.3.2
I/O commands defined in this specification					
0b	111 10b	01b	79h	Zone Management Send	3.4.3
0b	111 10b	10b	7Ah	Zone Management Receive	3.4.2
0b	111 11b	01b	7Dh	Zone Append	3.4.1

Notes:
1. Opcodes not listed are defined in the NVM Express Base Specification and in the NVM Command Set Specification.
2. All Zoned Namespace Command Set Commands use the Namespace Identifier (NSID) field. The value FFFFFFFFh is not supported in this field unless footnote 4 in this figure indicates that a specific command does support that value.
3. Indicates the data transfer direction of the command. All options to the command shall transfer data as specified or transfer no data. All commands, including vendor specific commands, shall follow this convention: 00b = no data transfer; 01b = host to controller; 10b = controller to host; 11b = bidirectional.
4. This command may support the use of the Namespace Identifier (NSID) field set to FFFFFFFFh.

ZNS Specification, Section 3.2[291]

[291] Ex. Z, ZNS Specification, at 18.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[292]

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[293]

[292] Ex. BB, NVMe Base Specification, at 107.

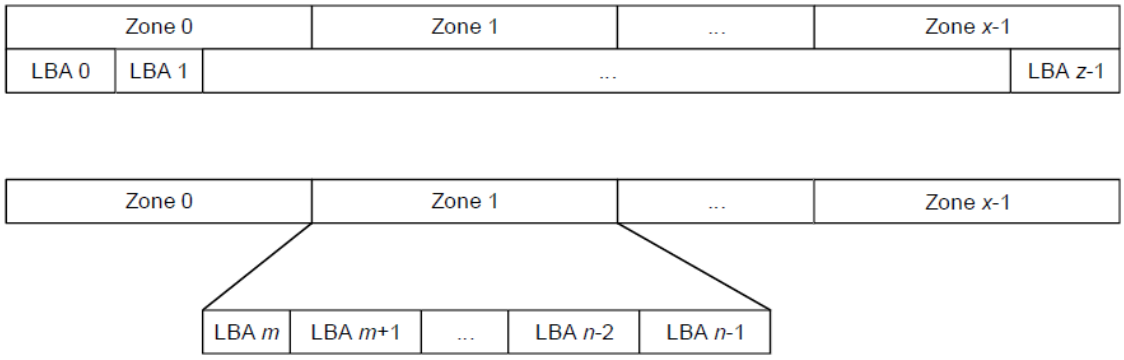
[293] Ex. AA, NVM Command Set Specification, at 54.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[294]

[294] Ex. Z, ZNS Specification, at 8.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.7[295]

1332. Samsung's Accused Instrumentalities contain circuitry to identify an addressed block device from the first address portion. '995 Patent, claim 1, element [1d].

1333. Samsung's Accused Instrumentalities are able to identify an addressed namespace (block device) from the Namespace Identifier or NSID (the first address portion). See:

[295] *Id*. at 20.

Figure 88: Common Command Format

Bytes	Description
07:04	**Namespace Identifier (NSID):** This field specifies the namespace that this command applies to. If the namespace identifier is not used for the command, then this field shall be cleared to 0h. The value FFFFFFFFh in this field is a broadcast value (refer to section 3.2.1.2), where the scope (e.g., the NVM subsystem, all attached namespaces, or all namespaces in the NVM subsystem) is dependent on the command. Refer to Figure 140 and Figure 391 for commands that support the use of the value FFFFFFFFh in this field. Specifying an inactive namespace identifier (refer to section 3.2.1.4) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Field in Command, unless otherwise specified. Specifying an invalid namespace identifier (refer to section 3.2.1.2) in a command that uses the namespace identifier shall cause the controller to abort the command with a status code of Invalid Namespace or Format, unless otherwise specified. If the namespace identifier is used for the command (refer to Figure 140), the value FFFFFFFFh is not supported for that command, and the host specifies a value of FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified. If the namespace identifier is not used for the command and the host specifies a value from 1h to FFFFFFFFh, then the controller should abort the command with a status code of Invalid Field in Command, unless otherwise specified.

NVMe Base Specification, Section 3.3.3.1[296]

1334. Samsung's Accused Instrumentalities contain circuitry to identify an addressed segment within the addressed block device from the second address portion. '995 Patent, claim 1, element [1e].

1335. Samsung's Accused Instrumentalities are able to identify an addressed zone (segment) within the addressed namespace from zone identifier information (second address portion). See:

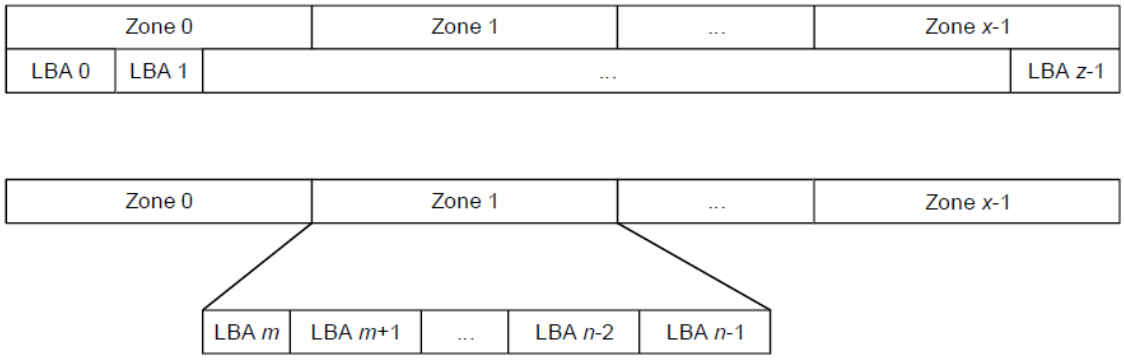
[296] Ex. BB, NVMe Base Specification, at 107.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[297]

[297] Ex. Z, ZNS Specification, at 8.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.5[298]

1336. Samsung's Accused Instrumentalities contain circuitry to identify an addressed storage location within the addressed segment from the third address portion. '995 Patent, claim 1, element [1f].

1337. Samsung's Accused Instrumentalities are able to identify an addressed storage location within the addressed zone from the LBA and/or alternatively a zone offset (the third address portion). See:

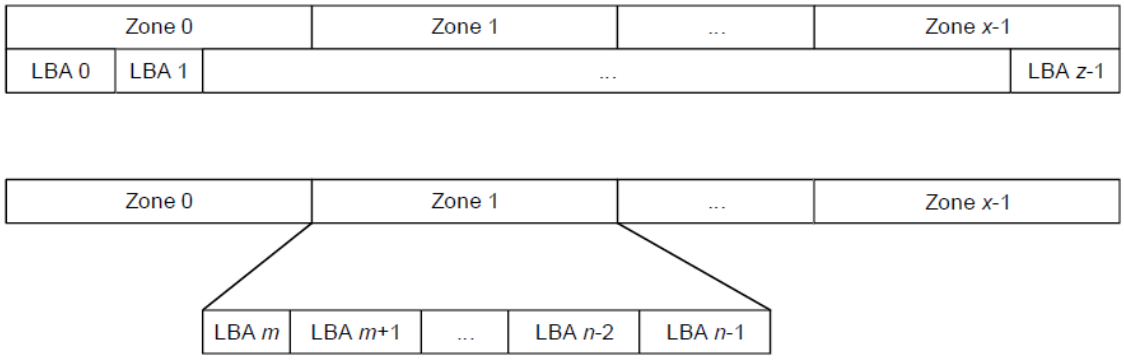
[298] *Id*. at 20.

2.1.1 Namespaces

A namespace is a collection of NVM and is defined in the NVM Express Base Specification and in the NVM Command Set Specification, as modified by this specification.

A zoned namespace is a namespace that is associated with the Zoned Namespace Command Set. A zoned namespace is divided into a set of equally-sized zones, which are contiguous non-overlapping ranges of logical block addresses. Figure 2 shows a zoned namespace with *x* zones and *z* LBAs where LBA 0 is the lowest LBA of zone 0, LBA *z*-1 is the highest LBA of zone *x*-1, and for Zone 1, *m* is the lowest LBA and *n*-1 is its highest LBA.

Figure 2: Zones in a Zoned Namespace



The Zoned Namespace Command Set is based on the NVM Command Set (refer to the NVM Command Set Specification).

Each zone has an associated Zone Descriptor that contains a set of attributes. A Zone Management Receive command may be used to retrieve one or more Zone Descriptors.

ZNS Specification, Section 2.1.1[299]

1.4.3.3 logical block

The smallest addressable data unit for Read and Write commands.

1.4.3.4 logical block address (LBA)

The address of a logical block, referred to commonly as LBA.

NVM Command Set Specification, Section 1.4[300]

[299] *Id*. at 8.

[300] Ex. AA, NVM Command Set Specification, at 10.

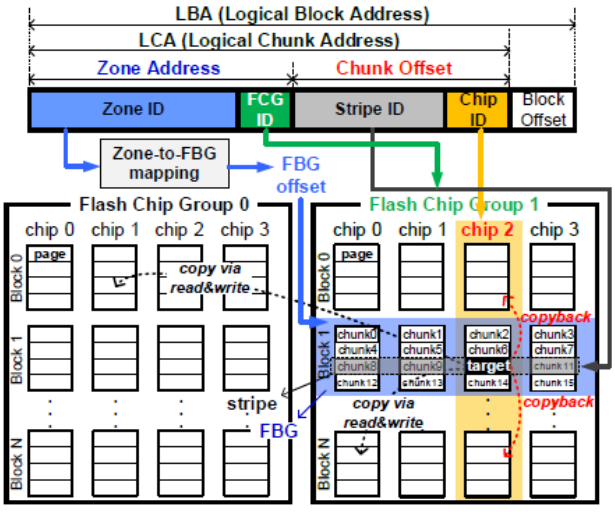


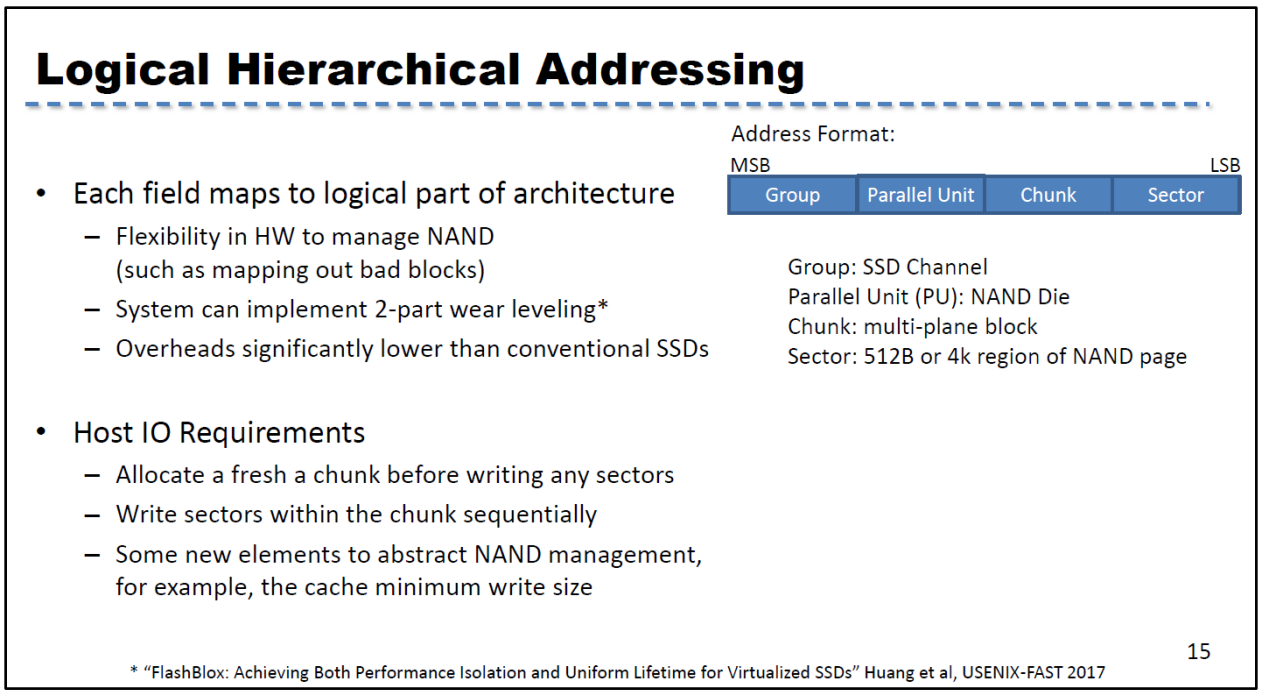
Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[301]

1338. Samsung's Accused Instrumentalities contain circuitry to use a division operation on at least one of the incoming address information or information derived therefrom to identify a logical erase unit within the addressed segment. '995 Patent, claim 1, element [1g].

1339. Samsung's Accused Instrumentalities are able to use a division operation on at least one of the incoming address information or information derived therefrom to identify a logical erase unit within the addressed zone. See:

[301] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.



Microsoft Denali Presentation[302]

[302] Ex. DD, Microsoft Denali Presentation, at 15.

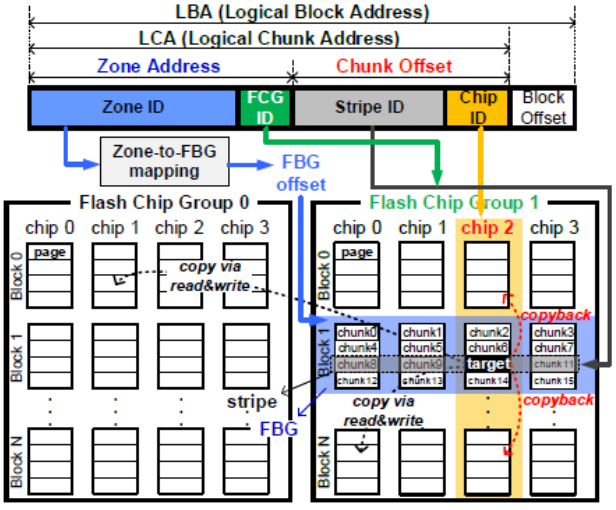
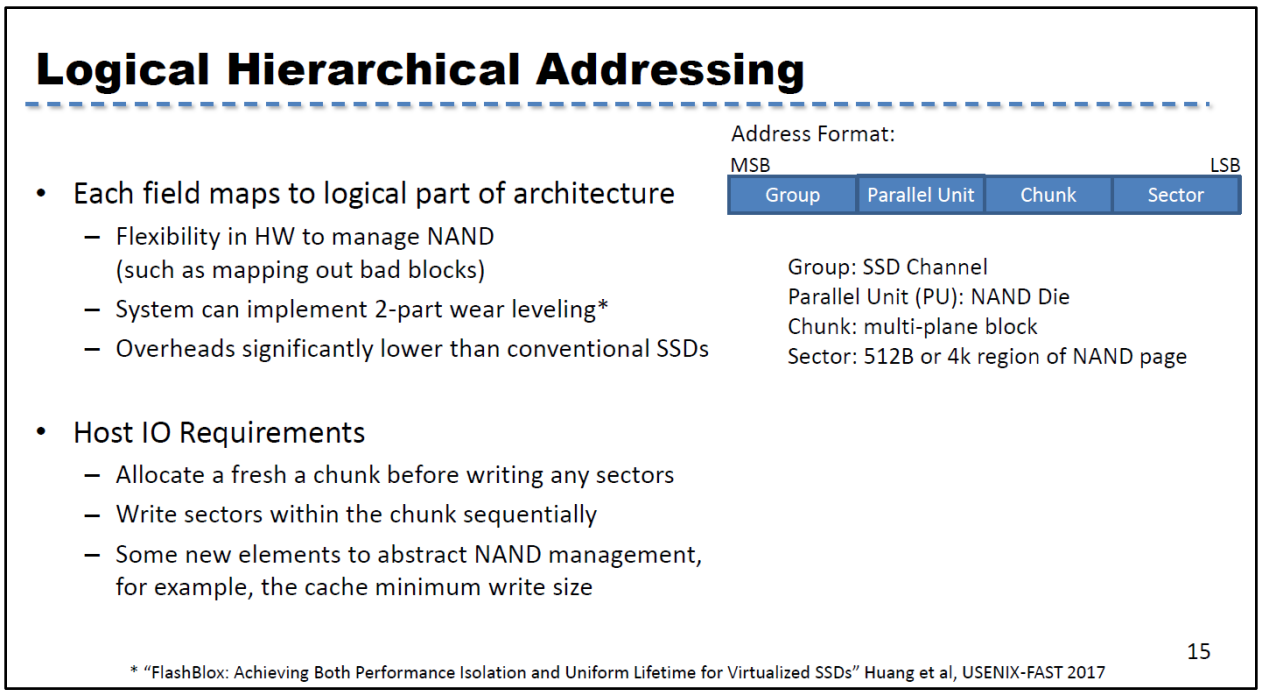


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[303]

1340. Samsung's Accused Instrumentalities contain circuitry to identify one of the physical erase units that corresponds to the identified logical erase unit. '995 Patent, claim 1, element [1h].

[303] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.



Microsoft Denali Presentation[304]

[304] Ex. DD, Microsoft Denali Presentation, at 15.

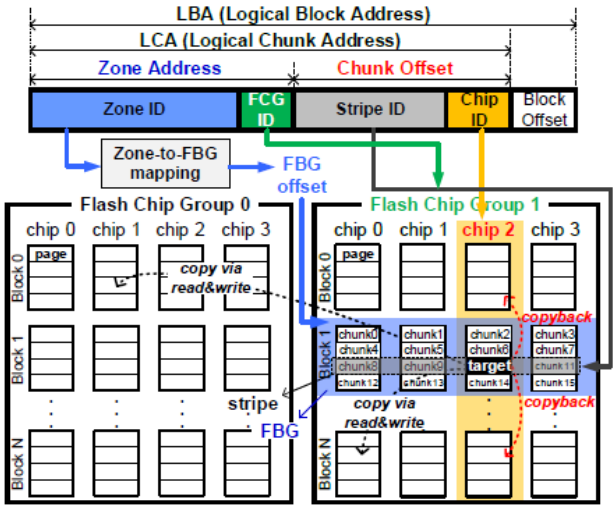


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[305]

1341. Samsung's Accused Instrumentalities contain circuitry to program associated data into the identified one of the physical erase units. '995 Patent, claim 1, element [1i].

1342. Samsung's Accused Instrumentalities are able to program data associated with the incoming Write command into the identified physical erase unit. See:

[305] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

3.3.7 Write command

The Write command operates as defined in the NVM Command Set Specification, with the additional requirements associated with the zone type of the specified zones that the command operates on (refer to section 2.1.1.2).

3.3.7.1 Command Completion

Command Completion is as defined in the NVM Command Set Specification, with the additional Zoned Namespace Command Set Command Specific status values that are defined in Figure 19.

Figure 19: Write – Command Specific Status Values

Value	Description
B8h	**Zone Boundary Error:** The command specifies logical blocks in more than one zone.
B9h	**Zone Is Full:** The accessed zone is in the ZSF:Full state.
BAh	**Zone Is Read Only:** The accessed zone is in the ZSRO:Read Only state.
BBh	**Zone Is Offline:** The accessed zone is in the ZSO:Offline state.
BCh	**Zone Invalid Write:** The write to a zone was not at the write pointer.
BDh	**Too Many Active Zones:** The controller does not allow additional active zones.
BEh	**Too Many Open Zones:** The controller does not allow additional open zones.

ZNS Specification, Section 3.3.5[306]

3.3.6 Write command

The Write command writes data and metadata, if applicable, to the I/O controller for the logical blocks indicated. The host may also specify protection information to include as part of the operation.

The command uses Command Dword 2, Command Dword 3, Command Dword 10, Command Dword 11, Command Dword 12, Command Dword 13, Command Dword 14, and Command Dword 15 fields. If the command uses PRPs for the data transfer, then the Metadata Pointer, PRP Entry 1, and PRP Entry 2 fields are used. If the command uses SGLs for the data transfer, then the Metadata SGL Segment Pointer and SGL Entry 1 fields are used.

Figure 67: Write – Data Pointer

Bits	Description
127:00	**Data Pointer (DPTR):** This field specifies the location of a data buffer where data is transferred from. Refer to the Common Command Format figure in the NVM Express Base Specification for the definition of this field.

[306] Ex. Z, ZNS Specification, at 20.

Figure 69: Write – Command Dword 10 and Command Dword 11

Bits	Description
63:00	**Starting LBA (SLBA):** This field indicates the 64-bit address of the first logical block to be written as part of the operation. Command Dword 10 contains bits 31:00; Command Dword 11 contains bits 63:32.

NVM Command Set Specification, Section 3.3.6[307]

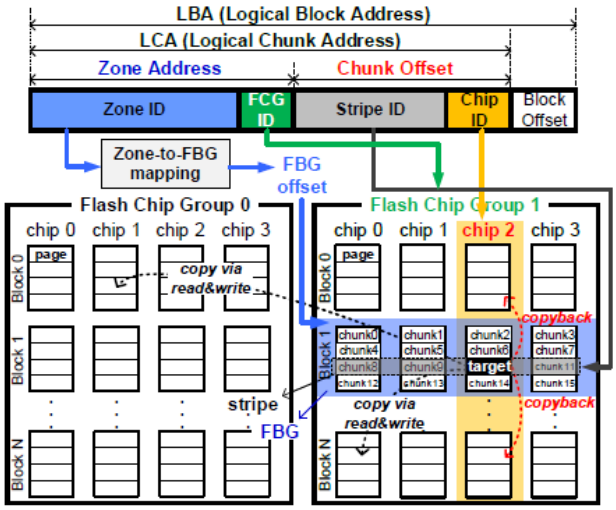


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[308]

1343. Samsung's Accused Instrumentalities satisfy the claim element "wherein the storage device also comprises circuitry to." '995 Patent, claim 1, element [1j]. Samsung's Accused Instrumentalities contain circuitry to perform elements [1k] through [1m]. See below.

1344. Samsung's Accused Instrumentalities contain circuitry to maintain metadata for

[307] Ex. AA, NVM Command Set Specification, at 53–54.

[308] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

each physical erase unit which contains data, the metadata representing a time since data was programmed in the respective erase unit. '995 Patent, claim 1, element [1k].

1345. Samsung's Accused Instrumentalities are able to maintain metadata for each physical block which contains data. On information and belief, such metadata includes metadata representing a time since data was programmed in the respective block. See:



Samsung Over-Provisioning White Paper[309]

[309] Ex. V, Samsung Over-Provisioning White Paper, at 1.

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-
194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[310]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

[310] *Id*. at 6.

Samsung SSD White Paper[311]

> **2. Why is OP important?**
>
> OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

Samsung Application Note[312]

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[313]

> **Reliability.** Endurance, in drive writes per day (DWPD), goes to 0.26 for the new device from 0.18 for the previous generation. Power-off data retention is increased from 1 month to 3 months.

Samsung Tech Blog[314]

---

[311] Ex. I, Samsung SSD White Paper, at 19.

[312] Ex. CC, Samsung Application Note, at 2.

[313] Ex. Z, ZNS Specification, at 41.

[314] Ex. HH, SAMSUNG, *Next-Generation QLC V-NAND Increases Data Center Profitability*, available at: https://semiconductor.samsung.com/news-events/tech-blog/next-generation-qlc-v-nand-increases-data-center-profitability/.

2.2.3. *Reliability limitation*

A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (*bitflips*) may happen in the stored data. These errors are due to the NAND flash memory cells' density inside the chip, and to the high voltages applied to the cells during the write and erase operations [RIC 14b, CHE 07]. Ultimately, this lack of reliability means that without a specific management the data written in the flash memory and that read some time afterwards will not necessarily be read the same. Moreover, in order to reduce the errors due to reliability problems of a NAND flash, during a write operation on the flash it is usually advisable to program the pages within a same block in a sequential way. This is a strong constraint for MLC chips, and a (strong) suggestion for SLC chips [GRU 09].

Data retention is a metric that indicates the time during which a cell is capable of memorizing its content after it was programmed. The retention time of flash memory cells is reported to be between 5 and 10 years [CHE 07]. These numbers should be weighted carefully, as the retention capability of a flash memory cell depends on several parameters, in particular on the number of erase operations performed during the programming. The numbers above suppose a number of erase operations equal to zero. The retention can drop to 1 year for chips at the end of their life [GRU 09].

Flash Memory Integration[315]

1346. Samsung's Accused Instrumentalities contain circuitry to copy data from a first one of the physical erase units for which the maintained metadata meets at least one criterion to a second one of the physical erase units. '995 Patent, claim 1, element [1l].

1347. Samsung's Accused Instrumentalities are able to copy data from a first one of the physical blocks for which the maintained metadata meets a threshold to a second one of the physical blocks. See:

[315] Ex. W, FLASH MEMORY INTEGRATION, at 26.



Samsung Over-Provisioning White Paper[316]

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-

[316] Ex. V, Samsung Over-Provisioning White Paper, at 1.

194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[317]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

[317] *Id*. at 6.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

Samsung SSD White Paper[318]

2. Why is OP important?

OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

[318] Ex. I, Samsung SSD White Paper, at 19.

Samsung Application Note[319]

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[320]

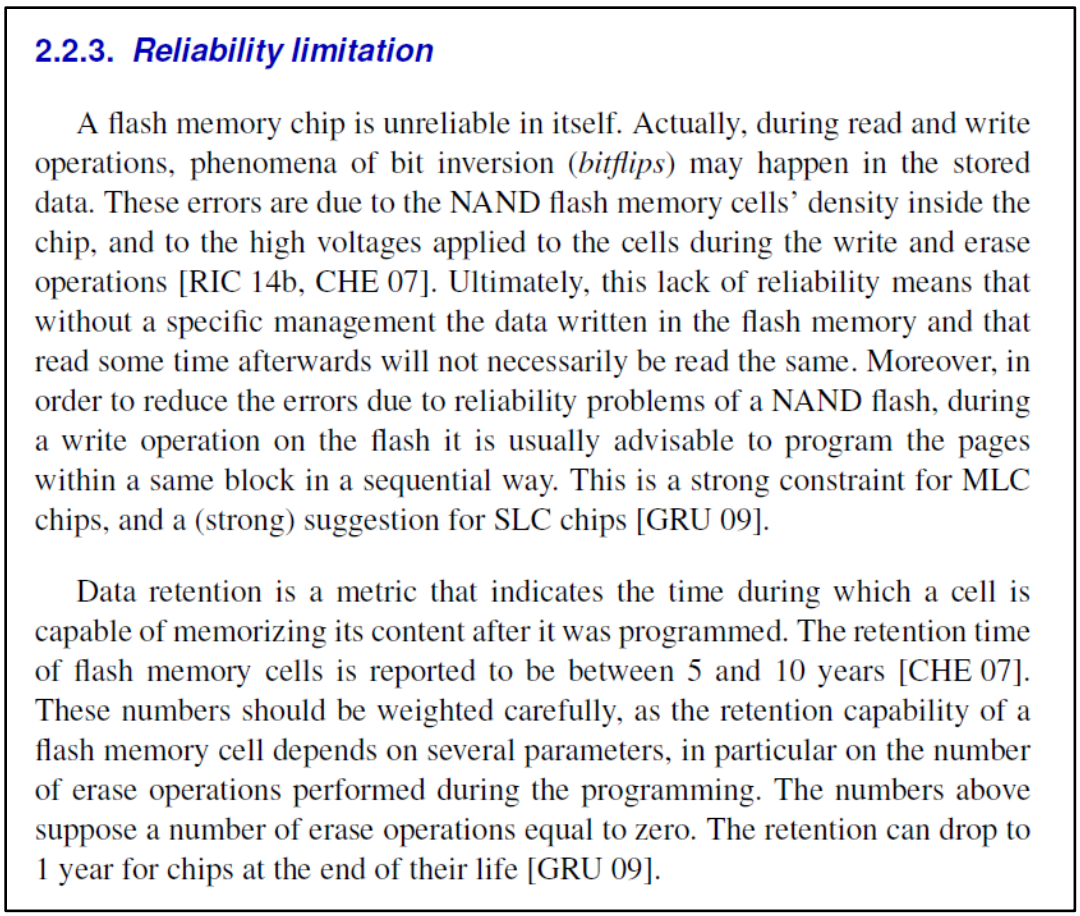
> **Reliability.** Endurance, in drive writes per day (DWPD), goes to 0.26 for the new device from 0.18 for the previous generation. Power-off data retention is increased from 1 month to 3 months.

Samsung Tech Blog[321]

[319] Ex. CC, Samsung Application Note, at 2.

[320] Ex. Z, ZNS Specification, at 41.

[321] Ex. HH, *Next-Generation QLC V-NAND Increases Data Center Profitability*.



2.2.3. *Reliability limitation*

A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (*bitflips*) may happen in the stored data. These errors are due to the NAND flash memory cells' density inside the chip, and to the high voltages applied to the cells during the write and erase operations [RIC 14b, CHE 07]. Ultimately, this lack of reliability means that without a specific management the data written in the flash memory and that read some time afterwards will not necessarily be read the same. Moreover, in order to reduce the errors due to reliability problems of a NAND flash, during a write operation on the flash it is usually advisable to program the pages within a same block in a sequential way. This is a strong constraint for MLC chips, and a (strong) suggestion for SLC chips [GRU 09].

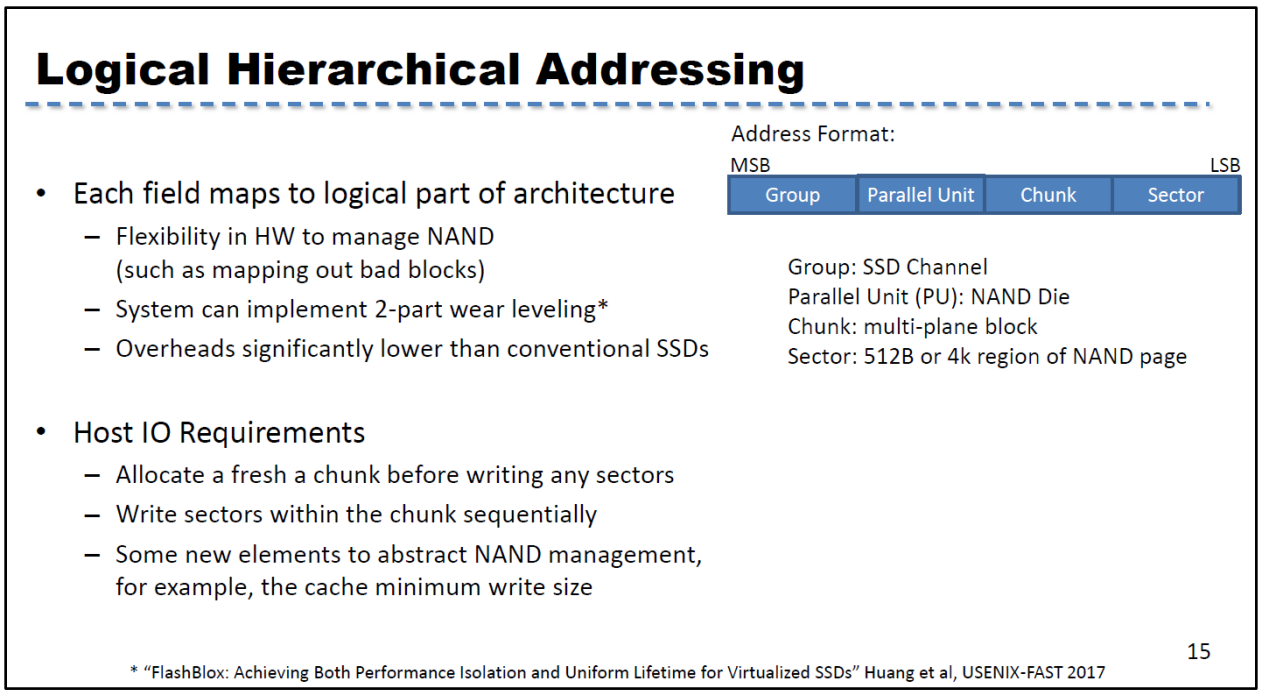
Data retention is a metric that indicates the time during which a cell is capable of memorizing its content after it was programmed. The retention time of flash memory cells is reported to be between 5 and 10 years [CHE 07]. These numbers should be weighted carefully, as the retention capability of a flash memory cell depends on several parameters, in particular on the number of erase operations performed during the programming. The numbers above suppose a number of erase operations equal to zero. The retention can drop to 1 year for chips at the end of their life [GRU 09].

Flash Memory Integration[322]

1348. Samsung's Accused Instrumentalities contain circuitry to remap a correspondence of a logical erase unit which corresponds to the first one of the physical erase units to instead correspond to the second one of the physical erase units. '995 Patent, claim 1, element [1m].

1349. On information and belief, after the data is moved from the first physical block to the second physical block, Samsung's Accused Instrumentalities remap the correspondence of the logical block so that it instead corresponds to the second physical block. See:

[322] Ex. W, FLASH MEMORY INTEGRATION, at 26.



Microsoft Denali Presentation[323]

[323] Ex. DD, Microsoft Denali Presentation, at 15.

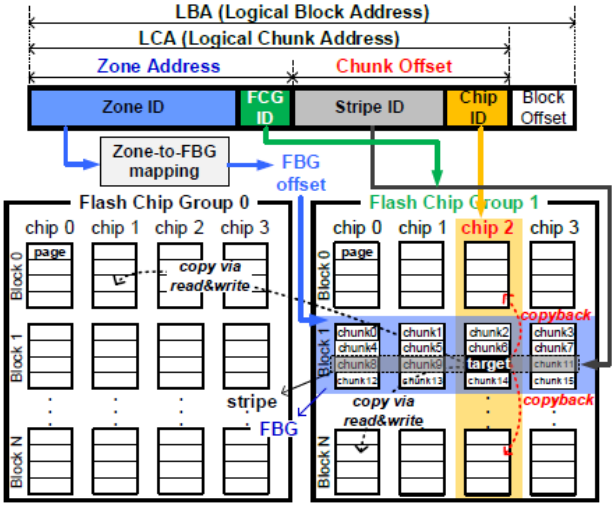


Figure 1: An example of zone and chunk mapping. With the zone address, the second FBG in the FCG 1 is selected. With the chunk offset, the third stripe in the selected FBG and the third flash page (chip 2) in the stripe are targeted.

Samsung Research Paper[324]



Samsung Over-Provisioning White Paper[325]

[324] Ex. EE, *ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction*, at 149.

[325] Ex. V, Samsung Over-Provisioning White Paper, at 1.

ID	Attribute name	Status Flag	Threshold (%)
5	Reallocated Sector Count	110011	10
9	Power-on Hours	110010	-
12	Power-on Count	110010	-
177	Wear Leveling Count	010011	5
179	Used Reserved Block Count (total)	010011	10
180	Unused Reserved Block Count (total)	010011	10
181	Program Fail Count (total)	110010	10
182	Erase Fail Count (total)	110010	10
183	Runtime Bad Count (total)	010011	10
184	End to End Error data path Error count	110011	97
187	Uncorrectable Error Count	110010	-
190	Airflow Temperature	110010	-
194	Temperature	100010	-
195	ECC Error Rate	011010	-
197	Pending Sector Count	110010	-
199	CRC Error Count	111110	-
202	SSD Mode Status	110011	10
235	POR Recovery Count	010010	-
241	Total LBAs Written	110010	-
242	Total LBAs Read	110010	-
243	SATA Downshift Control	110010	-
244	Thermal Throttle Status	110010	-
245	Timed Workload Media Wear	110010	-
246	Timed Workload Host Read / Write Ratio	110010	-
247	Timed Workload Timer	110010	-
251	NAND Writes	110010	-

Samsung Over-Provisioning White Paper[326]

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implantation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

[326] *Id*. at 6.

Samsung SSD White Paper[327]

> **2. Why is OP important?**
>
> OP has a direct effect on the SSD's random performance as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks (GC, wear-leveling, bad block management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the Write Amplification Factor (WAF) by ensuring there's room to work, which improves the SSD's lifetime.

Samsung Application Note[328]

> **5.4 Reset Zone Recommended**
>
> A controller that schedules an internal operation (e.g., background operation on the non-volatile media) on a zone that is in the ZSF:Full state may notify host software to initiate a zone reset operation (refer

ZNS Specification, Section 5.4[329]

> **Reliability.** Endurance, in drive writes per day (DWPD), goes to 0.26 for the new device from 0.18 for the previous generation. Power-off data retention is increased from 1 month to 3 months.

Samsung Tech Blog[330]

---

[327] Ex. I, Samsung SSD White Paper, at 19.

[328] Ex. CC, Samsung Over-Provisioning White Paper, at 2.

[329] Ex. Z, ZNS Specification, at 41.

[330] Ex. HH, *Next-Generation QLC V-NAND Increases Data Center Profitability*.

2.2.3. *Reliability limitation*

A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (*bitflips*) may happen in the stored data. These errors are due to the NAND flash memory cells' density inside the chip, and to the high voltages applied to the cells during the write and erase operations [RIC 14b, CHE 07]. Ultimately, this lack of reliability means that without a specific management the data written in the flash memory and that read some time afterwards will not necessarily be read the same. Moreover, in order to reduce the errors due to reliability problems of a NAND flash, during a write operation on the flash it is usually advisable to program the pages within a same block in a sequential way. This is a strong constraint for MLC chips, and a (strong) suggestion for SLC chips [GRU 09].

Data retention is a metric that indicates the time during which a cell is capable of memorizing its content after it was programmed. The retention time of flash memory cells is reported to be between 5 and 10 years [CHE 07]. These numbers should be weighted carefully, as the retention capability of a flash memory cell depends on several parameters, in particular on the number of erase operations performed during the programming. The numbers above suppose a number of erase operations equal to zero. The retention can drop to 1 year for chips at the end of their life [GRU 09].

Flash Memory Integration[331]

1350. Accordingly, Samsung's Accused Instrumentalities contain each and every element in claim 1 of the '995 Patent.

1351. SEC's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1352. SEC's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

[331] Ex. W, Flash Memory Integration, at 26.

1353. SEC's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1354. SEC directly infringes the Asserted Patent by using the Accused Instrumentalities.

1355. SEC directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1356. SEC directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1357. SEC directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1358. SEC directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1359. SEC directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1360. SEA's customers directly infringe the Asserted Patent by using the Accused Instrumentalities in the United States.

1361. SEA's customers directly infringe the Asserted Patent by testing the Accused Instrumentalities in the United States.

1362. SEA's customers directly infringe the Asserted Patent by making the Accused Instrumentalities in the United States.

1363. SEA directly infringes the Asserted Patent by using the Accused Instrumentalities in the United States.

1364. SEA directly infringes the Asserted Patent by testing the Accused Instrumentalities in the United States.

1365. SEA directly infringes the Asserted Patent by making the Accused Instrumentalities in the United States.

1366. SEA directly infringes the Asserted Patent by selling the Accused Instrumentalities in the United States.

1367. SEA directly infringes the Asserted Patent by offering to sell the Accused Instrumentalities in the United States.

1368. SEA directly infringes the Asserted Patent by importing the Accused Instrumentalities into the United States.

1369. SEC induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities.

1370. SEC induces infringement by SEA by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1371. SEC induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities.

1372. SEC contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the

Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1373. SEC contributes to infringement by SEA by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-

infringing use.

1374. SEC contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEC has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1375. SEC has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1376. SEC took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1377. SEA induces Samsung's customers' infringement by offering for sale, selling, encouraging the making of, encouraging the use of, instructing on the making of, instructing on the use of, promoting, and servicing the Accused Instrumentalities in the United States.

1378. SEA induces infringement by other Samsung subsidiaries by financing, supporting, encouraging, directing, controlling, and coordinating the research, design, development, building, manufacturing, sale, distribution, maintenance, testing, advertising, marketing, and/or operation of the Accused Instrumentalities in the United States.

1379. SEA contributes to infringement by Samsung's customers by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1380. SEA contributes to infringement by other Samsung subsidiaries by offering to sell, selling, distributing, or otherwise providing, in the United States, components that constitute a material part of the inventions in the Asserted Patents and components that are used to practice one or more processes or methods claimed in the Asserted Patents. These components for the Accused Instrumentalities include memory controllers, NAND flash, firmware, software, APIs and

libraries associated with the firmware or software, software or hardware tools used to install or facilitate the installation of the Accused Instrumentalities or their components, software or hardware tools used to assemble or facilitate the assembly of the Accused Instrumentalities or their components, and software or hardware tools used to configure or facilitate the configuration of the Accused Instrumentalities or their components. In offering to sell, selling, distributing, or otherwise providing these components, SEA has known or has had reason to believe that these components are especially made or especially adapted for use in an infringement of the Asserted Patents and that these components are not a staple article or commodity of commerce suitable for substantial non-infringing use.

1381. SEA has notice of the Asserted Patent and Radian's infringement allegations at least as of the filing of this Complaint.

1382. SEA took the above actions intending to cause infringing acts by others, and/or it willfully blinded itself as to the existence of the Asserted Patent, the Accused Instrumentalities' infringement thereof, and the infringing nature of others' actions.

1383. The acts that Radian alleges to give rise to infringement liability, as described above, are continuing.

1384. The acts that Radian alleges to give rise to infringement liability, as described above, will continue unless enjoined.

1385. SEC's infringement of the Asserted Patent is ongoing.

1386. SEC will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1387. SEC's ongoing infringement of the Asserted Patent is willful.

1388. SEC's ongoing infringement of the Asserted Patent is egregious.

1389. SEA's infringement of the Asserted Patent is ongoing.

1390. SEA will not stop its infringement of the Asserted Patent absent a Court order restraining it from future infringement.

1391. SEA's ongoing infringement of the Asserted Patent is willful.

1392. SEA's ongoing infringement of the Asserted Patent is egregious.

1393. Samsung has no acceptable non-infringing alternatives to the Asserted Patent.

1394. Hard disk drives (or "HDDs") are not an acceptable non-infringing alternative to the Asserted Patent.

1395. Conventional SSDs with a FTL are not an acceptable non-infringing alternative to the Asserted Patent.

1396. Flexible Data Placement (or "FDP") SSDs are not an acceptable non-infringing alternative to the Asserted Patent.

1397. SEC's actions have caused damage to Radian. Radian is entitled to recover from SEC the damages sustained by Radian as a result of SEC's wrongful acts in an amount subject to proof at trial.

1398. SEA's actions have caused damage to Radian. Radian is entitled to recover from SEA the damages sustained by Radian as a result of SEA's wrongful acts in an amount subject to proof at trial.

PRAYER FOR RELIEF

WHEREFORE, Radian prays for the following relief:

(a) A finding that Defendants infringe the patents-in-suit;

(b) A finding that Defendants are liable for infringement of the patents-in-suit;

(c) A judgment awarding damages against Defendants for their infringement of the patents-in-suit;

(d) A finding and judgment that the patents-in-suit are valid and enforceable;

(e) Preliminary and permanent injunctions enjoining Defendants and Defendants' agents, officers, employees, assigns and others acting in concert with them from committing any further act of infringement, whether direct or indirect, of the patents-in-suit;

(f) An award of damages adequate to compensate Plaintiff for any infringement by the Defendants or their agents that is not restrained prior to judgment;

(g) An award of pre- and post-judgment interest until the judgment is satisfied;

(h) A finding that this is an egregious case and an award of enhanced damages and other relief provided by 35 U.S.C. § 284;

(i) A finding that this is an exceptional case and an award of Plaintiff's attorneys' fees pursuant and other relief provided by 35 U.S.C. § 285;

(j) An award to Plaintiff of its costs; and

(k) Any further relief as the Court may deem just and equitable.

JURY DEMAND

Plaintiff demands a trial by jury on all issues so triable.

DATED: December 23, 2024

Respectfully submitted,

/s/ Jason D. Cassady

Jason D. Cassady
Texas State Bar No. 24045625
Email: jcassady@caldwellcc.com
Bradley W. Caldwell
Texas State Bar No. 24040630
Email: bcaldwell@caldwellcc.com
John Austin Curry
Texas State Bar No. 24059636
acurry@caldwellcc.com
Hamad M. Hamad
Texas State Bar No. 24061268
Email: hhamad@caldwellcc.com
Adrienne R. Dellinger
Texas Bar No. 24116275
Email: adellinger@caldwellcc.com
Alexander J. Gras
Texas Bar No. 24125252
Email: agras@caldwellcc.com
CALDWELL CASSADY & CURRY P.C.
2121 N Pearl Street, Suite 1200
Dallas, TX 75201
Telephone: (214) 888-4848

Andrea L. Fair
Texas Bar No. 24078488
Email: andrea@millerfairhenry.com
MILLER FAIR HENRY, PLLC
1507 Bill Owens Parkway
Longview, Texas 75604
Telephone: (903) 757-6400
Facsimile: (903) 757-2323

Attorneys for Radian Memory Systems LLC