


Exhibit 3
U.S. Patent No. 11,681,614

Sources:

- Samsung, Samsung Solid State Drive White Paper, available at: https://download.semiconductor.samsung.com/resources/white-paper/Samsung_SSD_White_Paper.pdf (herein “Samsung SSD White Paper”).
- Samsung, Samsung PM1733 NVMe SSD Product Brief, available at: <https://download.semiconductor.samsung.com/resources/brochure/PM1733%20NVMe%20SSD.pdf> (herein “Samsung PM1733 Product Brief”).
- 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 (herein “Samsung Over-Provisioning White Paper”).
- Jalil Boukhobza et al., Flash Memory Integration (ISTE Press Ltd 2017) (herein “Boukhobza”).
- 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 (herein “Bae”).
- Zoned Storage, SSDs with NVMe Zoned Namespace (ZNS) Support, available at: <https://zonedstorage.io/docs/introduction/zns> (herein “Zoned Storage Introduction”).
- NVM Express, Zoned Namespace Command Set Specification, Revision 1.1d, (herein “ZNS Specification”).
- NVM Express, NVM Express Base Specification, Revision 2.0d, (herein “NVMe Base Specification”).
- 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 (herein “Samsung Application Note”).
- Microsoft, Denali the Next Generation High Density Storage Interface, March 2018 OCP Summit (herein “Microsoft Denali Presentation”).
- Kyuhwa Han et al., ZNS+: Advanced Zoned Namespace Interface for Supporting In-Storage Zone Compaction, 15th USENIX Symposium on Operating Systems Design and Implementation (herein “Han”).

- 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/> (herein “EETimes V-NAND Analysis”).
- 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> (herein “Anand V-NAND Analysis”).
- 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/> (herein “Samsung Tech Blog”).
- Matias Bjørling et al., USENIX, Getting Zoned Storage out of the Friend Zone, available at <https://www.usenix.org/publications/loginonline/getting-zoned-storage-out-friend-zone> (herein “Bjørling”).
- Samsung's U.S. Patent Application Publication No. 2021/0349781 (herein “Samsung's '781 Patent Application”).
- Samsung's U.S. Patent Application Publication No. 2024/0319890 (herein “Samsung's '890 Patent Application”).
- Samsung's Chinese Patent Application Publication No. CN117,055,816A (herein “Samsung's '816 Patent Application”).
- Samsung's U.S. Patent No. 11,334,250 (herein “Samsung's '250 Patent”).

No.	Claim Language	Samsung's Accused Instrumentalities
<i>Ipre</i>	A storage device comprising:	To the extent the preamble is considered a limitation, Samsung's Accused Instrumentalities comprise a storage device.
<i>Ia</i>	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:	Samsung's Accused Instrumentalities comprise 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. Samsung's Accused Instrumentalities comprise flash memory. For example, Samsung's PM1733 comprises flash memory. See:



Source: Samsung PM1733 Product Brief (annotated).

Samsung's Accused Instrumentalities are able to divide memory into zones. 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. Source: <i>Id.</i> (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.

Source: ZNS Specification, at 8.

Each zone is a subdivision comprising a set of blocks (erase units). See:

What Is ZNS?

- Zoned namespaces (ZNS): Emerging storage interface
 - Divide logical address space into multiple **zones**
 - In general, each zone is mapped to **one or more flash blocks**



Source: Bae, at 4.

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.

Source: Zoned Storage Introduction.

Each block (erase unit) comprises host-accessible 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

Source: 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).

Figure 2: Garbage Collection



Source: 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


		Source: Boukhobza, at 18 (highlighted). – <i>Erase operation</i>, which is applied at the level of an entire block’s granularity. Therefore, the erase affects <u>all the pages contained in a given block</u>. Source: <i>Id.</i> at 22 (highlighted).
<i>Ib</i>	is associated with a logical address range,	Samsung’s Accused Instrumentalities comprise zones wherein each zone is associated with a logical address range. 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.

Source: ZNS Specification, at 8.

		<table><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></table> Source: <i>Id.</i> at 28. <table><tr><td>63:00</td><td>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.</td></tr></table> Source: <i>Id.</i> 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.	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.
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.							
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.							
<i>Ic</i>	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,	Samsung's Accused Instrumentalities comprise 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. Each zone has a zone size (zone capacity). See: <table><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></table> Source: ZNS Specification, at 28.	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.		
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.							

Samsung's Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action (a management command) from a host that results in the storage device performing an erasure of one or more blocks (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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: ZNS Specification, 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:Implicitly 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. Source: <i>Id.</i> at 30–31. 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. Source: ZNS Specification, at 42. MB: Yes, George is right. What has been standardized today is that, for ZNS, you need to write sequentially. If you want to write to a specific page again, you need a “zone reset”—which erases the zone's contents and reverts the zone to an empty state—and then you can write to that page again. But inherently, with an SSD, you
--	--	---


		Source: Bjørling.
<i>1d</i>	is in a full state when there are no host accessible pages available for writing, and	Samsung's Accused Instrumentalities comprise zones wherein each zone is in a full state when there are no host accessible pages available for writing. A zone is in a full state when there are no host accessible pages available for writing. See: Figure 7: Zone State Machine <pre> graph LR subgraph Active_Resources [Active Resources] ZSE((ZSE:Empty)) ZSF((ZSF:Full)) ZSC((ZSC:Closed)) ZSE <--> ZSF ZSC --> ZSE ZSC --> ZSF end subgraph Open_Resources [Open Resources] ZSIO((ZSIO: Implicitly Opened)) ZSEO((ZSEO: Explicitly Opened)) ZSIO --> ZSEO end ZSIO --> ZSC ZSC --> ZSIO ZSIO --> ZSRO((ZSRO:Read Only)) ZSRO --> ZSO((ZSO:Offline)) ZSEO --> ZSO </pre> The diagram illustrates the Zone State Machine. It features two main containers: 'Active Resources' and 'Open Resources'. 'Active Resources' contains three states: ZSE:Empty, ZSF:Full, and ZSC:Closed. ZSE:Empty and ZSF:Full are connected by a bidirectional arrow. ZSC:Closed has arrows pointing to both ZSE:Empty and ZSF:Full. 'Open Resources' contains two states: ZSIO: Implicitly Opened and ZSEO: Explicitly Opened, connected by a downward arrow. ZSIO: Implicitly Opened has an arrow pointing to ZSC:Closed. ZSC:Closed has an arrow pointing back to ZSIO: Implicitly Opened. ZSIO: Implicitly Opened has an arrow pointing to ZSRO:Read Only, which then points to ZSO:Offline. ZSEO: Explicitly Opened also has an arrow pointing to ZSO:Offline. Source: ZNS Specification, at 12.

		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). Source: <i>Id.</i> 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). Source: <i>Id.</i>
<i>Ie</i>	is in an empty state when all host accessible pages are available for writing; and	Samsung's Accused Instrumentalities comprise zones wherein each zone is in an empty state when all host accessible pages are available for writing. A zone is in an empty state when all host accessible pages are available for writing. See:



Source: ZNS Specification, at 12.

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.



Source: *Id.* at 10.

		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'; Source: <i>Id.</i> at 11. 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. Source: <i>Id.</i> at 12. 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. Source: <i>Id.</i> 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. Source: <i>Id.</i> 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. Source: <i>Id.</i> at 14. All logical blocks in a zone shall be marked as deallocated when the zone is in the ZSE:Empty state. Source: <i>Id.</i> at 15.
--	--	---

<i>If</i>	logic operable to:	Samsung's Accused Instrumentalities comprise "logic operable to." Samsung's Accused Instrumentalities comprise logic operable to perform elements [1g] through [1j]. See below.
<i>1g</i>	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	Samsung's Accused Instrumentalities comprise 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. Samsung's Accused Instrumentalities are able to receive a Zone Management Receive command specifying a Report Zones action (query command) from the host and, in response, execute the command and return information about the 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. Source: <i>Id.</i> at 24. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', Samsung's Accused Instrumentalities are able to return information to list all zones, as well as the zone state, zone capacity (zone size), zone start LBA ("ZSLBA"), and write pointer of all zones. See:

	respective zones include host accessible pages available for writing;	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. Source: <i>Id.</i>	
--	--	--	--

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.

Source: *Id.* at 26.

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
		Eh	Full
	Fh	Offline	
	3:0	Reserved	

Figure 37: Zone Descriptor Data Structure



Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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	

Source: *Id.* at 27–28.

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. Source: <i>Id.</i> 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  Source: <i>Id.</i> at 10.
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	Samsung’s Accused Instrumentalities comprise 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. 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).

	the set of erase units; and	Source: ZNS Specification, at 20. Samsung's Accused Instrumentalities are also able to receive and execute a Zone Append command from the host. The Zone Append command is a command to write data. 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. Source: <i>Id.</i> at 21. Samsung's Accused Instrumentalities are able to perform a multi-plane write such that the data is written concurrently to two or more blocks (erase units) of a specified 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 mm² while the 48L V-NAND die measures 99.8 mm², 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/mm². The memory density of leading-edge 2D planar NAND devices such as Toshiba 15 nm TLC NAND is 1.28 Gb/mm², (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 mm² to 68.7 mm² 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. Source: EETimes V-NAND Analysis (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

Source: Anand V-NAND Analysis.




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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7


The diagram illustrates a system architecture with four columns (11a, 11b, 11c, 11d) and four rows (Z1, Z2, Z3, Z4). Each column contains a vertical stack of units. Units in columns 11a, 11b, and 11c are labeled EU1-EU4, EU5-EU8, and EU9-EU12 respectively. Units in column 11d are labeled RA1-RA5, RA6-RA8, and RA9-RA12 respectively. Each unit is connected to a corresponding STRP (Storage/Transfer/Processing) unit. The diagram illustrates a data flow from left to right, with units in columns 11a, 11b, and 11c feeding into units in column 11d. The units in column 11d are then connected to STRP units, which are further connected to a common bus (CNS).

		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域 (Zone) 的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a), 当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b), 当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
<i>li</i>	update the information representing the extent to which the specified zone includes host accessible pages available for writing,	Samsung’s Accused Instrumentalities comprise logic operable to update the information representing the extent to which the specified zone includes host accessible pages available for writing. Samsung’s Accused Instrumentalities are able to update the write pointer (information representing the extent to which the specified 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. Source: ZNS Specification, at 9.

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.

Source: *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



Source: *Id.*


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:Implicitly 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. Source: <i>Id.</i> at 30–31.	
<i>Ij</i>	wherein the logic comprises at least one of hardware logic or instructions stored on non-transitory machine-readable media that, when executed,	Samsung's Accused Instrumentalities comprise said logic comprising at least one of hardware logic or instructions stored on non-transitory machine-readable media that, when executed, control operation of hardware circuitry. See, for example:	

control
operation of
hardware
circuitry.




Source: Samsung PM1733 Product Brief (annotated).

		09 Why Integration Matters What Samsung's Vertical Integration Means to You The three most crucial components of any SSD are the NAND flash memory, the controller, and the firmware. The first does the critical task of storing your precious data; the second two work together to accomplish the complex and equally essential tasks of managing data storage and maintaining the performance and lifespan of the SSD. 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. Redundantly stored in NAND flash, the firmware may be updated manually to improve or extend functionality when the SSD manufacturer releases an update. 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. Thus, the firmware, which drives the controller, is arguably the most important component of a good SSD because it delivers the largest difference among branded SSDs in terms of performance, reliability, and endurance (lifespan). Developing good firmware requires precise engineering skills and painstaking integration among the NAND flash, controller, DRAM and other SSD components. Intimate knowledge of NAND characteristics, semiconductor process technologies, and controller characteristics must go hand-in-hand. Samsung represents one of the few manufacturers that designs all of these components in-house; and, among those who do, Samsung enjoys the most experience (with over 6 years as the number one supplier in the preinstalled SSD market). Source: Samsung SSD White Paper, at 30 (highlighted).
2	The storage device of claim 1, wherein when the specified zone has no host accessible pages	Samsung's Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung's Accused Instrumentalities comprise the storage device operable to update a state of the specified zone to the full state when the specified zone has no host accessible pages available for writing after writing data to the set of erase units of the specified zone using the multi-plane write command.

available for writing after writing data to the set of erase units of the specified zone using the multi-plane write command, the storage device is operable to update a state of the specified zone to the full state.

Samsung's Accused Instrumentalities are able to update the state of a zone to Full when the zone has no available host accessible pages after writing to the set of blocks (erase units) of the specified zone using the multi-plane write command. See:



Source: ZNS Specification, at 12.

		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). 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). Source: <i>Id.</i> at 13.
3	The storage device of claim 1, wherein the logic is operable to write the data to the two or more erase units of the set of erase units of the specified zone	Samsung's Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung's Accused Instrumentalities comprise the logic operable to write the data to the two or more erase units of the set of erase units of the specified zone using a common page address. Samsung's Accused Instrumentalities are able to write the data to the two or more blocks (erase units) of the specified zone using a common page address. See:

using a common page address.





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域(Zone)的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a),当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b),当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
4	The storage device of claim 1, wherein the zone size is a minimum memory size associated with the management command.	Samsung’s Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung’s Accused Instrumentalities comprise the storage device of claim 1 wherein the zone size is a minimum memory size associated with the management command. Samsung’s Accused Instrumentalities are able to receive and execute the Zone Management Send command specifying a Reset Zone action (management command) which specifies at least a zone and is associated with the memory size of at least a 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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: ZNS Specification, 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:Implicitly 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.
		Source: <i>Id.</i> at 30–31.
5	The storage device of claim 1, wherein the management command from the host results in the logic issuing erase commands to the set of erase	Samsung's Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung's Accused Instrumentalities comprise the storage device of claim 1 wherein the management command from the host results in the logic issuing erase commands to the set of erase units of the specified zone that are on two or more different planes. Samsung's Accused Instrumentalities are able to receive and execute the Zone Management Send command specifying a Reset Zone action (management command) from the host which results in the logic issuing erase commands to the set of blocks (erase units) of the specified zone that are on two or more different planes. See:

units of the specified zone that are on two or more different planes.

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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: ZNS Specification, 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:

- 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:Implicitly 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.
		Source: <i>Id.</i> at 30–31.

		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 mm² while the 48L V-NAND die measures 99.8 mm², 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/mm². The memory density of leading-edge 2D planar NAND devices such as Toshiba 15 nm TLC NAND is 1.28 Gb/mm², (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 mm² to 68.7 mm² 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. Source: EETimes V-NAND Analysis (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

Source: Anand V-NAND Analysis.





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <u>Translation:</u> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域(Zone)的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a),当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b),当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <u>Translation:</u>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
6	The storage device of claim 1, wherein the information about the zones is stored on non-volatile memory of the storage device.	Samsung’s Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung’s Accused Instrumentalities comprise the storage device of claim 1 wherein the information about the zones is stored on non-volatile memory of the storage device. The information about the zone has to be maintained by Samsung’s Accused Instrumentalities when power is off, and therefore must be stored on non-volatile memory of the storage device.
7	The storage device of claim 1, wherein the information about the zones is stored on a die that includes the logic.	Samsung’s Accused Instrumentalities comprise the storage device of claim 1. See claim 1. Samsung’s Accused Instrumentalities comprise the storage device of claim 1 wherein the information about the zones is stored on a die that includes the logic. Samsung’s Accused Instrumentalities store the information about the zones on a die that includes the logic.

<i>9pre</i>	The storage device of claim 1, wherein the storage device is operable to:	Samsung's Accused Instrumentalities comprise the storage device of claim 1. See claim 1.
<i>9a</i>	generate metadata indicating an amount of time since the specified zone was written to, and	Samsung's Accused Instrumentalities comprise the storage device operable to generate metadata indicating an amount of time since the specified zone was written to. On information and belief, Samsung's Accused Instrumentalities generate metadata including metadata indicating an amount of time since a specified zone was written to, including metadata such as age of data or data retention period. See:





FIG. 4

		[0067] The storage media 125 may have a retention period (which may depend on the usage history of the storage media 125, and, as such, may vary within the storage media 125); data that has been stored longer than the retention period (i.e., data having an age exceeding the retention period) may become unreliable and may be said to have expired. Data may be stored in the storage media 125 using an error correcting code, which may be a block code. When data is read from the storage media 125, a quantity of raw data, referred to as a code block, may be read from the storage media 125, and an attempt to decode it may be made. If the attempt fails, additional attempts (read retrials) may be made. With use, a portion, e.g., a block, of the storage media 125 may degrade to the point that the retention period becomes unacceptably short, and the block may be classified as a “bad block”. To avoid allowing this circumstance to render the entire storage media 125 inoperable, reserve space, referred to as “bad block management reserve space” may be present (e.g., included in each flash memory die or in each flash memory plane), and the controller 120, or another controller internal to the flash memory die or to the flash memory plane may begin to use a block in the reserve and cease to use the bad block.	
--	--	--	--

	[0089] FIG. 4 is a schematic data layout diagram of a storage device 110 in accordance with example embodiments of the disclosure. A block lifespan table 405 lists the retention period of each block in the storage device 110, and a mapping table 410 lists mappings from logical block address to physical block addresses. When the storage device 110 transitions to the sustainable read-only mode, the storage device 110 copies data from the faulty used space 415 (the physical space affected by the fault condition) to the rescue space 420. Data in the normal used space 425 is handled as it would be if the storage device 110 were in normal operating mode (instead of being in the sustainable read-only mode). Because the fault condition results in a loss of physical space, there is a corresponding reduction in the size of the used LBA space 430, with some of the previously available LBA space becoming unused LBA space 435. Firmware 440, executed by the controller 120, may cause the data to be moved, as described above, and the tables to be updated.	
--	---	--

[0090] FIG. 5A is a flow chart for a method of transitioning to the sustainable read-only mode, in accordance with example embodiments of the disclosure. At 505, the host 105 sends/receives data from storage device; at 510, the host 105 determines if an error has occurred in the storage device; at 520, the host 105 determines if an error happens with the storage device; at 525, the storage device 110 performs an internal diagnosis and transitions to the sustainable read-only mode; at 530, an FPGA may optionally monitor the nature of valid data on the storage device 110 (e.g., based on data type, or age); at 535, the storage device 110 relocates valid data to the rescue space when the age of the data approaches the retention period; and at 540, the storage device 110 may optionally transmit information about the status and contents of the rescue space to the host 105, which flushes the data. FIG. 5B is a flow chart showing details of

Source: Samsung's '781 Patent Application at Fig. 4, [0067], [0089]–[0090].



Source: 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	-

Source: *Id.* at 6.

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 implementation 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. Source: 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. Source: Samsung Application Note, at 2. 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. Source: Samsung Tech Blog.
--	--	---

		2.2.3. Reliability limitation A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (<i>bitflips</i>) 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]. Source: Boukhobza, at 26.	
9b	in response to comparing the metadata to a threshold,	Samsung's Accused Instrumentalities comprise storage device operable to, in response to comparing the metadata to a threshold, provide an indication that the storage device recommends that the host send a second management command directed to the specified zone.	

provide an indication that the storage device recommends that the host send a second management command directed to the specified zone.	Samsung's Accused Instrumentalities are able to detect the need for management based on comparing the metadata to a threshold, provide a "Reset Zone Recommended" indication that the storage device recommends that the host send a Zone Management Send command specifying a Reset Zone action (second management command) directed to the 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 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'. Source: ZNS Specification, at 41–42.
--	---

		Figure 37: Zone Descriptor Data Structure <table><tr><th>Bytes</th><th>Description</th></tr><tr><td rowspan="6">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></table> Source: <i>Id.</i> at 28. See also claim element [9a].	Bytes	Description	02	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>	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.
Bytes	Description																	
02	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>	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.				
	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.																	
10pr e	A storage device comprising:	To the extent the preamble is considered a limitation, Samsung’s Accused Instrumentalities comprise a storage device.																
10a	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	Samsung’s Accused Instrumentalities comprise 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. Samsung’s Accused Instrumentalities comprise flash memory. For example, Samsung’s PM1733 comprises flash memory. See:																

pages, wherein
each zone:



Source: Samsung PM1733 Product Brief (annotated).

Samsung's Accused Instrumentalities are able to divide memory into zones. 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. Source: <i>Id.</i> (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.

Source: ZNS Specification, at 8.

Each zone is a subdivision comprising a set of blocks (erase units). See:

What Is ZNS?

- Zoned namespaces (ZNS): Emerging storage interface
 - Divide logical address space into multiple **zones**
 - In general, each zone is mapped to **one or more flash blocks**



Source: Bae, at 4.

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.

Source: Zoned Storage Introduction.

Each block (erase unit) comprises host-accessible 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

Source: 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).

Figure 2: Garbage Collection



Source: 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


		Source: Boukhobza, at 18 (highlighted). – <i>Erase operation</i>, which is applied at the level of an entire block’s granularity. Therefore, the erase affects <u>all the pages contained in a given block</u>. Source: <i>Id.</i> at 22 (highlighted).
10b	is associated with a logical address range,	Samsung’s Accused Instrumentalities comprise zones wherein each zone is associated with a logical address range. 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.

Source: ZNS Specification, at 8.

		<table border="1"> <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> </table> Source: <i>Id.</i> at 28. <table border="1"> <tr> <td>63:00</td><td> 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. </td></tr> </table> Source: <i>Id.</i> 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.	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.
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.							
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.							
10c	has a zone size associated with a management command from a host,	Samsung's Accused Instrumentalities comprise zones wherein each zone has a zone size associated with a management command from a host. Each zone has a zone size (zone capacity). See: <table border="1"> <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> </table> Source: ZNS Specification at 28.	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.		
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.							

Samsung's Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action (a management command) from a host that results in the erasure of blocks (erase units) within the zone which is associated with the zone size. 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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: *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:


- 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:Implicitly 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. Source: <i>Id.</i> at 30–31. 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. Source: ZNS Specification, at 42. MB: Yes, George is right. What has been standardized today is that, for ZNS, you need to write sequentially. If you want to write to a specific page again, you need a “zone reset”—which erases the zone's contents and reverts the zone to an empty state—and then you can write to that page again. But inherently, with an SSD, you
--	--	---

		Source: Björling.
10d	transitions to a full state in response to no longer having host accessible pages available for writing, and	Samsung's Accused Instrumentalities comprise zones wherein each zone transitions to a full state in response to no longer having host accessible pages available for writing. A zone transitions to a full state when there are no longer host accessible pages available for writing. See: Figure 7: Zone State Machine <pre> graph LR subgraph Active_Resources [Active Resources] ZSE((ZSE:Empty)) ZSF((ZSF:Full)) ZSC((ZSC:Closed)) ZSE <--> ZSF ZSC -.-> ZSE end subgraph Open_Resources [Open Resources] ZSIO((ZSIO: Implicitly Opened)) ZSEO((ZSEO: Explicitly Opened)) ZSIO --> ZSEO end ZSIO --> ZSE ZSC <--> ZSEO ZSRO((ZSRO:Read Only)) ZSO((ZSO:Offline)) ZSRO --> ZSO ZSIO --> ZSRO ZSEO --> ZSO </pre> The diagram illustrates the Zone State Machine. It features two main containers: 'Active Resources' and 'Open Resources'. 'Active Resources' contains three states: ZSE:Empty, ZSF:Full, and ZSC:Closed. ZSE:Empty and ZSF:Full are connected by a bidirectional arrow. A dashed arrow points from ZSC:Closed to ZSE:Empty. 'Open Resources' contains two states: ZSIO: Implicitly Opened and ZSEO: Explicitly Opened, connected by a downward arrow. Transitions include: ZSIO: Implicitly Opened to ZSE:Empty; ZSC:Closed to ZSEO: Explicitly Opened; ZSIO: Implicitly Opened to ZSRO:Read Only; and ZSRO:Read Only to ZSO:Offline. ZSEO: Explicitly Opened also transitions to ZSO:Offline. Source: ZNS Specification, at 12.


		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). Source: <i>Id.</i> 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). Source: <i>Id.</i>
10e	in response to the storage device receiving the management command associated therewith, transitions to an empty state such that all	Samsung's Accused Instrumentalities comprise zones wherein each zone, in response to the storage device receiving the management command associated therewith, transitions to an empty state such that all host accessible pages are available for writing. A zone is in an empty state when all host accessible pages are available for writing. See:

host accessible
pages are
available for
writing; and



Source: ZNS Specification at 12.

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.



Source: *Id.* at 10.

		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'; Source: <i>Id.</i> at 11. 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. Source: <i>Id.</i> at 12. 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. Source: <i>Id.</i> 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. Source: <i>Id.</i> 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. Source: <i>Id.</i> at 14. All logical blocks in a zone shall be marked as deallocated when the zone is in the ZSE:Empty state. Source: <i>Id.</i> at 15.
--	--	---

		A zone transitions to the empty state in response to a Zone Management Send command specifying a Reset Zone action (management command) associated therewith. 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:Implicitly 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. Source: <i>Id.</i> at 30–31.
10f	logic operable to:	Samsung's Accused Instrumentalities comprise "logic operable to."

		Samsung's Accused Instrumentalities comprise logic operable to perform elements [10g] through [10p]. See below.
10g	return information about zones stored by the storage device in response to a query command from the host, the information about the zones comprising:	Samsung's Accused Instrumentalities comprise logic operable to return information about zones stored by the storage device in response to a query command from the host. Samsung's Accused Instrumentalities are able to receive a Zone Management Receive command specifying a Report Zones action (query command) from the host and, in response, execute the command and return information about zones stored by the storage device. 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. Source: ZNS Specification, at 24.
10h	a list of the zones accessible by the host,	Samsung's Accused Instrumentalities comprise logic operable to return a list of the zones accessible by the host. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', Samsung's Accused Instrumentalities are able to return a list of the zones accessible by the host. 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.	
		Source: ZNS Specification, at 24.	

		Figure 35: Report Zones Data Structure <table><tr><th>Bytes</th><th>Description</th></tr><tr><td>07:00</td><td>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.</td></tr><tr><td>63:08</td><td>Reserved</td></tr><tr><td>127:64</td><td>Zone Descriptor 0: Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).</td></tr><tr><td>191:128</td><td>Zone Descriptor 1: Contains the Zone Descriptor for the second zone reported, if any.</td></tr><tr><td>...</td><td>...</td></tr><tr><td>$((n+1)*64)+63:(n+1)*64$</td><td>Zone Descriptor n: Contains the Zone Descriptor for the last zone reported, if any.</td></tr></table> Source: <i>Id.</i> at 26.	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.
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.															
10i	a respective logical address associated with a zone of the zones stored,	Samsung’s Accused Instrumentalities comprise logic operable to return a respective logical address associated with a zone of the zones stored. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to ‘0h’, Samsung’s Accused Instrumentalities are able to return the zone start logical block address (“ZSLBA”) associated with each zone. See:														

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
	Eh	Full	
Fh	Offline		
3:0	Reserved		

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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.												

		<table> <tr> <td></td><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> <tr> <td>07:04</td><td colspan="2">Reserved</td></tr> <tr> <td>15:08</td><td colspan="2"> 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 colspan="2"> 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 colspan="2"> 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 colspan="2">Reserved</td></tr> </table> Source: ZNS Specification, at 27–28.		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	
	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																			
10j	the zone size,	Samsung's Accused Instrumentalities comprise logic operable to return the zone size. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', Samsung's Accused Instrumentalities are able to return the zone capacity (zone size). See:																		

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
	Eh	Full	
Fh	Offline		
3:0	Reserved		

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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.												

		<table> <tr> <td></td><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> <tr> <td>07:04</td><td colspan="2">Reserved</td></tr> <tr> <td>15:08</td><td colspan="2"> 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 colspan="2"> 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 colspan="2"> 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 colspan="2">Reserved</td></tr> </table> Source: ZNS Specification, at 27–28.		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	
	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																			
10k	information indicating, for respective zones in which there are no host accessible pages available for writing the respective zone, that the respective	Samsung's Accused Instrumentalities comprise logic operable to return information indicating, for respective zones in which there are no host accessible pages available for writing the respective zone, that the respective zones are in the full state. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', Samsung's Accused Instrumentalities are able to return the zone state, including, for respective zones in which there are no host accessible pages available for writing, that the respective zones are in the full state. See:																		

zones are in the full state, and

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
		Eh	Full
	Fh	Offline	
3:0	Reserved		

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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.												

		<table border="1"> <tr> <td></td><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> <tr> <td>07:04</td><td colspan="2">Reserved</td></tr> <tr> <td>15:08</td><td colspan="2"> 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 colspan="2"> 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 colspan="2"> 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 colspan="2">Reserved</td></tr> </table> Source: ZNS Specification, at 27–28.		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	
	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																			
10l	information representing, for respective zones, an extent to which the respective zones include host accessible pages	Samsung's Accused Instrumentalities comprise logic operable to return information representing, for respective zones, an extent to which the respective zones include host accessible pages available for writing. In response to a Zone Management Receive command specifying a Report Zones action with the reporting option set to '0h', Samsung's Accused Instrumentalities are able to return a write pointer, which indicates an extent to which a zone includes host accessible pages available for writing. See:																		

available for writing;

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
		Eh	Full
	Fh	Offline	
3:0	Reserved		


Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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	

Source: ZNS Specification, at 27–28.

		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. Source: <i>Id.</i> 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  Source: <i>Id.</i> at 10.
10m	write data to the set of erase units of a specified zone using a common page address;	Samsung's Accused Instrumentalities comprise logic operable to write data to the set of erase units of a specified zone using a common page address. Samsung's Accused Instrumentalities are able to write data to the set of blocks (erase units) of a specified zone using a common page address. See:





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域(Zone)的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a),当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b),当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
10n	update the information representing the extent to which the specified zone includes host accessible pages available for writing; and	Samsung’s Accused Instrumentalities comprise logic operable to update the information representing the extent to which the specified zone includes host accessible pages available for writing. Samsung’s Accused Instrumentalities are able to update the write pointer (information representing the extent to which the specified 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. Source: ZNS Specification, at 9. 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.

Source: *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



Source: *Id.*


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:Implicitly 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.	
		Source: <i>Id.</i> at 30–31.	
10o	when the specified zone has no host accessible pages available for writing the zone after writing data to the set of erase units of the specified zone, update the	Samsung's Accused Instrumentalities comprise logic operable to, when the specified zone has no host accessible pages available for writing the zone after writing data to the set of erase units of the specified zone, update the information indicating that the specified zone is in the full state. Samsung's Accused Instrumentalities are able to update the zone state to Full when the specified zone has no host accessible pages available for writing after writing data to that zone. See:	

information indicating that the specified zone is in the full state,




Source: ZNS Specification, at 12.

Transition ZSIO:ZSF: The zone shall transition from the ZSIO:Implicitly Opened state to the ZSF:Full state:

- as a result of successful completion of a Zone Management Send command with a Zone Send Action of Finish Zone;
- as a result of a write operation that writes one or more logical blocks that causes the zone to reach its writeable zone capacity;
- due to a Zone Active Excursion (refer to section 5.6); and
- as a result of the zoned namespace becoming write protected (refer to the Namespace Write Protection section in the NVM Express Base Specification).

		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). Source: <i>Id.</i> at 13.
10p	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.	Samsung's Accused Instrumentalities comprise said logic comprising at least one of hardware logic or instructions stored on non-transitory machine-readable media that, when executed, control operation of hardware circuitry. See, for example:




Source: Samsung PM1733 Product Brief (annotated).

		09 Why Integration Matters What Samsung's Vertical Integration Means to You The three most crucial components of any SSD are the NAND flash memory, the controller, and the firmware. The first does the critical task of storing your precious data; the second two work together to accomplish the complex and equally essential tasks of managing data storage and maintaining the performance and lifespan of the SSD. 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. Redundantly stored in NAND flash, the firmware may be updated manually to improve or extend functionality when the SSD manufacturer releases an update. 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. Thus, the firmware, which drives the controller, is arguably the most important component of a good SSD because it delivers the largest difference among branded SSDs in terms of performance, reliability, and endurance (lifespan). Developing good firmware requires precise engineering skills and painstaking integration among the NAND flash, controller, DRAM and other SSD components. Intimate knowledge of NAND characteristics, semiconductor process technologies, and controller characteristics must go hand-in-hand. Samsung represents one of the few manufacturers that designs all of these components in-house; and, among those who do, Samsung enjoys the most experience (with over 6 years as the number one supplier in the preinstalled SSD market). Source: Samsung SSD White Paper, at 30 (highlighted).
11	The storage device of claim 10, wherein the logic is operable to write the data to the set of	Samsung's Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung's Accused Instrumentalities comprise the logic to write the data to the set of erase units of the specified zone using a multi-plane write command such that the data is concurrently written to two or more erase units of the set of erase units.

erase units of the specified zone using a multi-plane write command such that the data is concurrently written to two or more erase units of the set of erase units.


Samsung's Accused Instrumentalities are able to write the data to the set of blocks (erase units) of the specified zone using a multi-plane write command such that the data is concurrently written to two or more blocks (erase units) of the zone. 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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

“Large Zone

- Spanning multiple dies
- High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域 (Zone) 的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a), 当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b), 当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
12	The storage device of claim 10, wherein the management command from the host results in the logic issuing at least one erase command to erase one or more erase units of the set of erase units of the specified zone.	Samsung’s Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung’s Accused Instrumentalities comprise the storage device of claim 10 wherein the management command from the host results in the logic issuing at least one erase command to erase one or more erase units of the set of erase units of the specified zone. Samsung’s Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action (management command) from the host that results in the storage device the logic issuing at least one erase command to erase 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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: ZNS Specification, 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:Implicitly 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. Source: <i>Id.</i> at 30–31. 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. Source: ZNS Specification, at 42. MB: Yes, George is right. What has been standardized today is that, for ZNS, you need to write sequentially. If you want to write to a specific page again, you need a “zone reset”—which erases the zone's contents and reverts the zone to an empty state—and then you can write to that page again. But inherently, with an SSD, you
--	--	---

		Source: Bjørling.																																							
13	The storage device of claim 10, wherein the management command from the host results in the logic issuing an erase command associated with erase units of the specified zone that are on two or more different planes.	Samsung’s Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung’s Accused Instrumentalities comprise the storage device of claim 10 wherein the management command from the host results in the logic issuing an erase command associated with erase units of the specified zone that are on two or more different planes. Samsung’s Accused Instrumentalities are able to receive and execute a Zone Management Send command specifying a Reset Zone action (management command) from the host which results in the logic issuing an erase command associated with blocks (erase units) of the specified zone that are on two or more different planes. See: Figure 39: Zone Management Send – Command Dword 13 <table><tr><th>Bits</th><th>Description</th></tr><tr><td>31:09</td><td>Reserved</td></tr><tr><td>08</td><td>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.</td></tr><tr><td rowspan="10">07:00</td><td>Zone Send Action (ZSA): Defines the zone action to be performed for Zone Management Send.</td></tr><tr><td><table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table></td></tr></table> Source: ZNS Specification, at 29.	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.	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
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.																																								
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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										
	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																																								

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:Implicitly 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.

Source: *Id.* at 30–31.

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.

	Source: ZNS Specification, at 42. MB: Yes, George is right. What has been standardized today is that, for ZNS, you need to write sequentially. If you want to write to a specific page again, you need a “zone reset”—which erases the zone's contents and reverts the zone to an empty state—and then you can write to that page again. But inherently, with an SSD, you Source: Bjørling.
--	--





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:


- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域(Zone)的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a),当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b),当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
14	The storage device of claim 13, wherein, as a result of issuing the erase command, the storage device is operable to detect a bad erase unit and mark the bad erase unit as unusable for storage of new write data.	Samsung’s Accused Instrumentalities comprise the storage device of claim 13. See claim 13. Samsung’s Accused Instrumentalities comprise the storage device operable to detect a bad erase unit and mark the bad erase unit as unusable for storage of new write data as a result of issuing the erase command. Samsung’s Accused Instrumentalities are able to detect a bad erase block (erase unit) and mark such bad erase block as unusable for storage of new write data as a result of issuing the erase command. 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. Source: 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, at 2. Logical Hierarchical Addressing MSB Group Parallel Unit Chunk Sector LSB Address Format: Group: SSD Channel Parallel Unit (PU): NAND Die Chunk: multi-plane block Sector: 512B or 4k region of NAND page - 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 Source: Microsoft Denali Presentation, at 15.
15pr e	The storage device of claim 13, wherein, the	Samsung's Accused Instrumentalities comprise the storage device of claim 13. See claim 13.

	storage device is operable to:	
15a	generate metadata indicating a count of erasures of the erase units, and	Samsung's Accused Instrumentalities comprise the storage device operable to generate metadata indicating a count of erasures of the erase units. Samsung's Accused Instrumentalities are able to generate metadata, including metadata indicating a count of erasures of the erase units. See: After a sub-block erase operation, the erase manage module 43a may check an erase status of an erased sub-block and/or a sub-block adjacent to the erased sub-block. For example, the erase manage module 43a may sense memory cells of the erased sub-block to determine whether specific parameters exceed a reference cycle value. The erase manage module 43a may read data of sub-block(s) adjacent to the erased sub-block to detect erase-inhibition efficiency. For example, the erase manage module 43a may detect bit error rate (BER) based on data read from an erased sub-block. The erase manage module 43a may acquire and monitor wear-leveling information (e.g., erase count) on the erased sub-block. In addition, the erase manage module 43a may read data of the erased sub-block to monitor a variation in threshold voltages of selected memory cells and/or a variation in the bit error rate (BER). The erase manage module 43a may also read data of an unselected sub-block to detect a variation in a threshold voltage. The memory controller 40 may perform various procedures for compensating for insufficient erasing of a selected sub-block based on erase status information detected by the erase manage module 43a. Source: Samsung's '250 Patent at 5:16–38.



Source: 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	-

Source: *Id.* at 6.

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 implementation 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. Source: 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. Source: Samsung Application Note, at 2.
15b	in response to comparing the metadata to a threshold, provide an indication that the storage device recommends that the host send a second management command directed to the specified zone.	Samsung's Accused Instrumentalities comprise the storage device operable to, in response to comparing the metadata to a threshold, provide an indication that the storage device recommends that the host send a second management command directed to the specified zone. Samsung's Accused Instrumentalities are able to detect the need for management based on comparing the metadata to a threshold, provide a "Reset Zone Recommended" indication that the storage device recommends that the host send a Zone Management Send command specifying a Reset Zone action (second management command) directed to the 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

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'.

Source: ZNS Specification, at 41–42.

Figure 37: Zone Descriptor Data Structure

Bytes	Description											
02	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>	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
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.											

		Source: <i>Id.</i> at 28. See also claim element [15a].												
16	The storage device of claim 10, further comprising a host interface that comprises buffers to receive the information about the zones and queue the information about the zones for transmission to the host.	Samsung’s Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung’s Accused Instrumentalities comprise a host interface that comprises buffers to receive the information about the zones and queue the information about the zones for transmission to the host. Samsung’s Accused Instrumentalities implement protocols in the PCIe and NVMe specifications that allow the SSD controller to communicate with the host (“host interface”). See: <table><tr><th colspan="2">Samsung PM1733 specifications</th></tr><tr><td>Form factor</td><td>U.2 / 2.5"</td></tr><tr><td>Capacity</td><td>1.92 TB, 3.84 TB, 7.68TB, and 15.36TB</td></tr><tr><td>Host interface</td><td>PCIe Gen 3/4 x4</td></tr><tr><td>Spec Compliance</td><td>NVMe spec rev.1.3 PCI Express base specification rev. 4.0</td></tr><tr><td>NAND flash memory</td><td>Samsung V-NAND®</td></tr></table> Source: Samsung PM1733 Product Brief (annotated). The host interface comprises buffers to receive the information about the zones and queue the information about the zones for transmission to the host. 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 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®													

1.5.1 Admin Queue

The Admin Queue is the Submission Queue and Completion Queue with identifier 0. The Admin Submission Queue and corresponding Admin Completion Queue are used to submit administrative commands and receive completions for those administrative commands, respectively.

The Admin Submission Queue is uniquely associated with the Admin Completion Queue.

Source: NVMe Base Specification, at 18.

1.5.12 command completion

A command is completed when the controller has completed processing the command, has updated status information in the completion queue entry, and has posted the completion queue entry to the associated Completion Queue.

1.5.13 command submission

For memory-based transport model (e.g. PCIe) implementations, a command is submitted when a Submission Queue Tail Doorbell write has completed that moves the Submission Queue Tail Pointer value past the Submission Queue slot in which the command was placed.

For message-based transport model (e.g. NVMe over Fabrics) implementations, a command is submitted when a host adds a capsule to a Submission Queue.

1.5.14 controller

A controller is the interface between a host and an NVM subsystem. There are three types of controllers:

- a) I/O controllers;
- b) Discovery controllers; and
- c) Administrative controllers.

A controller executes commands submitted by a host on a Submission Queue and posts a completion on a Completion Queue. All controllers implement one Admin Submission Queue and one Admin Completion Queue. Depending on the controller type, a controller may also implement one or more I/O Submission Queues and I/O Completion Queues. When PCI Express is used as the transport, then a controller is a PCI Express function.

Source: NVMe Base Specification, at 19–20.

2.1 Memory-Based Transport Model

In the memory-based model, Submission and Completion Queues are allocated in memory.

Host software creates queues, up to the maximum supported by the controller. Typically, the number of command queues created is based on the system configuration and anticipated workload. For example, on a four core processor based system, there may be a queue pair per core to avoid locking and ensure data structures are created in the appropriate processor core's cache. Figure 6 provides a graphical representation of the queue pair mechanism, showing a 1:1 mapping between Submission Queues and Completion Queues. Figure 7 shows an example where multiple I/O Submission Queues utilize the same I/O Completion Queue on Core B. Figure 6 and Figure 7 show that there is always a 1:1 mapping between the Admin Submission Queue and Admin Completion Queue.

Figure 6: Queue Pair Example, 1:1 Mapping





Figure 7: Queue Pair Example, $n:1$ Mapping



A Submission Queue (SQ) is a circular buffer with a fixed slot size that the host software uses to submit commands for execution by the controller. The host software updates the appropriate SQ Tail doorbell register when there are one to n new commands to execute. The previous SQ Tail value is overwritten in the controller when there is a new doorbell register write. The controller fetches SQ entries in order from the Submission Queue and may execute those commands in any order.


Each submission queue entry is a command. Commands are 64 bytes in size. The physical memory locations in memory to use for data transfers are specified using Physical Region Page (PRP) entries or Scatter Gather Lists (SGL). Each command may include two PRP entries or one Scatter Gather List segment. If more than two PRP entries are necessary to describe the data buffer, then a pointer to a PRP List that describes a list of PRP entries is provided. If more than one SGL segment is necessary to describe the data buffer, then the SGL segment provides a pointer to the next SGL segment.

A Completion Queue (CQ) is a circular buffer with a fixed slot size used to post status for completed commands. A completed command is uniquely identified by a combination of the associated SQ identifier and command identifier that is assigned by host software. In the memory-based transport model multiple Submission Queues may be associated with a single Completion Queue. A configuration with a single Completion Queue may be used where a single worker thread processes all command completions via one Completion Queue even when those commands originated from multiple Submission Queues. The CQ Head pointer is updated by host software after processing completion queue entries indicating the last free CQ slot. A Phase Tag (P) bit is defined in the completion queue entry to indicate whether an entry has been newly posted without the host consulting a register (refer to section 3.3.3.2.2). The Phase Tag bit enables the host to determine whether entries are new or not.

		Source: NVMe Base Specification, at 29–30.
17	The storage device of claim 10, wherein the information about the zones is stored on non-volatile memory of the storage device.	Samsung's Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung's Accused Instrumentalities comprise the storage device of claim 10 wherein the information about the zones is stored on non-volatile memory of the storage device. The information about the zone has to be maintained by Samsung's Accused Instrumentalities when power is off, and therefore must be stored on non-volatile memory of the storage device.
18	The storage device of claim 10, wherein the information about the zones is stored on a die that includes the logic.	Samsung's Accused Instrumentalities comprise the storage device of claim 10. See claim 10. Samsung's Accused Instrumentalities comprise the storage device of claim 10 wherein the information about the zones is stored on a die that includes the logic. Samsung's Accused Instrumentalities store the information about the zones on a die that includes the logic.
20 <i>pr e</i>	The storage device of claim 10, wherein the storage device is operable to:	Samsung's Accused Instrumentalities comprise the storage device of claim 10. See claim 10.
20 <i>a</i>	generate metadata indicating an	Samsung's Accused Instrumentalities comprise the storage device operable to generate metadata indicating an amount of time since the specified zone was written to.

amount of time since the specified zone was written to, and

On information and belief, Samsung's Accused Instrumentalities generate metadata including metadata indicating an amount of time since a specified zone was written to, including metadata such as age of data or data retention period. See:




		[0067] The storage media 125 may have a retention period (which may depend on the usage history of the storage media 125, and, as such, may vary within the storage media 125); data that has been stored longer than the retention period (i.e., data having an age exceeding the retention period) may become unreliable and may be said to have expired. Data may be stored in the storage media 125 using an error correcting code, which may be a block code. When data is read from the storage media 125, a quantity of raw data, referred to as a code block, may be read from the storage media 125, and an attempt to decode it may be made. If the attempt fails, additional attempts (read retrials) may be made. With use, a portion, e.g., a block, of the storage media 125 may degrade to the point that the retention period becomes unacceptably short, and the block may be classified as a “bad block”. To avoid allowing this circumstance to render the entire storage media 125 inoperable, reserve space, referred to as “bad block management reserve space” may be present (e.g., included in each flash memory die or in each flash memory plane), and the controller 120, or another controller internal to the flash memory die or to the flash memory plane may begin to use a block in the reserve and cease to use the bad block.	
--	--	--	--

	[0089] FIG. 4 is a schematic data layout diagram of a storage device 110 in accordance with example embodiments of the disclosure. A block lifespan table 405 lists the retention period of each block in the storage device 110, and a mapping table 410 lists mappings from logical block address to physical block addresses. When the storage device 110 transitions to the sustainable read-only mode, the storage device 110 copies data from the faulty used space 415 (the physical space affected by the fault condition) to the rescue space 420. Data in the normal used space 425 is handled as it would be if the storage device 110 were in normal operating mode (instead of being in the sustainable read-only mode). Because the fault condition results in a loss of physical space, there is a corresponding reduction in the size of the used LBA space 430, with some of the previously available LBA space becoming unused LBA space 435. Firmware 440, executed by the controller 120, may cause the data to be moved, as described above, and the tables to be updated.	
--	---	--

[0090] FIG. 5A is a flow chart for a method of transitioning to the sustainable read-only mode, in accordance with example embodiments of the disclosure. At 505, the host 105 sends/receives data from storage device; at 510, the host 105 determines if an error has occurred in the storage device; at 520, the host 105 determines if an error happens with the storage device; at 525, the storage device 110 performs an internal diagnosis and transitions to the sustainable read-only mode; at 530, an FPGA may optionally monitor the nature of valid data on the storage device 110 (e.g., based on data type, or age); at 535, the storage device 110 relocates valid data to the rescue space when the age of the data approaches the retention period; and at 540, the storage device 110 may optionally transmit information about the status and contents of the rescue space to the host 105, which flushes the data. FIG. 5B is a flow chart showing details of

Source: Samsung's '781 Patent Application at Fig. 4, [0067], [0089]–[0090].



Source: 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	-

Source: *Id.* at 6.

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 implementation 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. Source: 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. Source: Samsung Application Note, at 2. 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. Source: Samsung Tech Blog.
--	--	---

		2.2.3. Reliability limitation A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (<i>bitflips</i>) 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]. Source: Boukhobza, at 26.	
20b	in response to comparing the metadata to a threshold,	Samsung's Accused Instrumentalities comprise a storage device operable to, in response to comparing the metadata to a threshold, provide an indication that the storage device recommends that the host send a second management command directed to the specified zone.	

provide an indication that the storage device recommends that the host send a second management command directed to the specified zone.	Samsung's Accused Instrumentalities are able to detect the need for management based on comparing the metadata to a threshold, and provide a "Reset Zone Recommended" indication that the storage device recommends that the host send a Zone Management Send command specifying a Reset Zone action (second management command) directed to the 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 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'. Source: ZNS Specification, at 41–42.
--	---

pages, and the zones are each associated with a respective logical address range; and



Source: Samsung PM1733 Product Brief (annotated).

Samsung's Accused Instrumentalities are able to divide memory into zones. 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. Source: <i>Id.</i> (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.

Source: ZNS Specification, at 8.

Each zone is a subdivision comprising a set of blocks (erase units). See:

What Is ZNS?

- Zoned namespaces (ZNS): Emerging storage interface
 - Divide logical address space into multiple **zones**
 - In general, each zone is mapped to **one or more flash blocks**



Source: Bae, at 4.

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.

Source: Zoned Storage Introduction.

Each block (erase unit) comprises host-accessible 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

Source: 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).

Figure 2: Garbage Collection



Source: 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


	Source: Boukhobza, at 18 (highlighted). – <i>Erase operation</i>, which is applied at the level of an entire block’s granularity. Therefore, the erase affects <u>all the pages contained in a given block</u>. Source: <i>Id.</i> at 22 (highlighted). Each zone is associated with a range of LBAs (logical address range). 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.



Source: ZNS Specification, at 8.

		<table><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></table> Source: <i>Id.</i> at 28. <table><tr><td>63:00</td><td> 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. </td></tr></table> Source: <i>Id.</i> 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.	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.
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.							
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.							
21b	a memory controller in communication with the flash memory, the memory controller comprising:	Samsung's Accused Instrumentalities comprise a memory controller in communication with the flash memory. Samsung's Accused Instrumentalities contain a memory controller in communication with the flash memory. For example, Samsung's PM1733 contains a memory controller (annotated in red) in communication with the flash memory (annotated in blue). See:						

		 Source: Samsung PM1733 Product Brief (annotated).
21c	storage arranged to retain	Samsung's Accused Instrumentalities comprise storage arranged to retain information about the zones, the information about the zones comprising zone size, a logical address associated with a zone of the zones,

information about the zones, the information about the zones comprising zone size, a logical address associated with a zone of the zones, and information representing an extent to which the zone of the zones comprises host accessible pages available for writing; and	and information representing an extent to which the zone of the zones comprises host accessible pages available for writing. Samsung's Accused Instrumentalities are able to retain information about the zones, including the zone capacity (zone size), the zone start LBA ("ZSLBA"), and a write pointer. See: <table border="1" data-bbox="531 383 1751 837"> <tr> <td data-bbox="531 383 653 691">15:08</td><td data-bbox="653 383 1751 691"> 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 data-bbox="531 691 653 748">23:16</td><td data-bbox="653 691 1751 748"> 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 data-bbox="531 748 653 805">31:24</td><td data-bbox="653 748 1751 805"> 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 data-bbox="531 805 653 837">63:32</td><td data-bbox="653 805 1751 837">Reserved</td></tr> </table> Source: ZNS Specification at 28. The write pointer is information representing an extent to which the zone of the zones comprises host accessible pages available for writing. See:	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
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								

		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. Source: <i>Id.</i> 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  Source: <i>Id.</i> at 10.
21d	logic operable to cause the storage device to output the information about the zones in response to a query command from a host, write to the set	Samsung's Accused Instrumentalities comprise logic operable to cause the storage device to output the information about the zones in response to a query command from a host, write to the set of erase units of the zone of the zones using a common page address, and update the information representing the extent to which the zone comprises host accessible pages available for writing in response to writing to the set of erase units of the zone. Samsung's Accused Instrumentalities are able to receive a Zone Management Receive command specifying a Report Zones action (query command) from the host and, in response, execute the command and return information about the zones. See:

of erase units of the zone of the zones using a common page address, and update the information representing the extent to which the zone comprises host accessible pages available for writing in response to writing to the set of erase units of the zone.	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. Source: ZNS Specification, at 24. Samsung’s Accused Instrumentalities are able to, in response to the Zone Management Receive command specifying a Report Zones action, cause the storage device to output the information about the 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. Source: <i>Id.</i>
--	--

15:08	Zone Receive Action Specific:																																								
	Zone Receive Action	Description																																							
	Report Zones	Reporting Options:																																							
		<table border="1"> <tr> <th>Value</th> <th>Description</th> </tr> <tr> <td>0h</td> <td>List all zones.</td> </tr> <tr> <td>1h</td> <td>List the zones in the ZSE:Empty state.</td> </tr> <tr> <td>2h</td> <td>List the zones in the ZSIO:Implicitly Opened state.</td> </tr> <tr> <td>3h</td> <td>List the zones in the ZSEO:Explicitly Opened state.</td> </tr> <tr> <td>4h</td> <td>List the zones in the ZSC:Closed state.</td> </tr> <tr> <td>5h</td> <td>List the zones in the ZSF:Full state.</td> </tr> <tr> <td>6h</td> <td>List the zones in the ZSRO:Read Only state.</td> </tr> <tr> <td>7h</td> <td>List the zones in the ZSO:Offline state.</td> </tr> <tr> <td>8h</td> <td>Reserved</td> </tr> <tr> <td>9h</td> <td> 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'. </td> </tr> <tr> <td>10h</td> <td rowspan="3">Reserved</td> </tr> <tr> <td>11h</td> </tr> <tr> <td>3Fh</td> </tr> <tr> <td>All other values</td> <td>Reserved</td> </tr> <tr> <td>Extended Report Zones</td> <td></td> </tr> <tr> <td>All other values</td> <td>Reserved</td> </tr> </table> <table border="1"> <tr> <th>Bits</th> <th>Description</th> </tr> <tr> <td>00h</td> <td>Report Zones: Reports Zone Descriptor entries through the Report Zones data structure (refer to Figure 35).</td> </tr> <tr> <td>01h</td> <td>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.</td> </tr> <tr> <td>02h to FFh</td> <td>Reserved</td> </tr> </table> Source: <i>Id.</i> at 25–26.	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	Extended Report Zones		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
	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																																								
Extended Report Zones																																									
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																																								

Figure 35: Report Zones Data Structure <table border="1"> <thead> <tr> <th>Bytes</th><th>Description</th></tr> </thead> <tbody> <tr> <td>07:00</td><td> 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. </td></tr> <tr> <td>63:08</td><td>Reserved</td></tr> <tr> <td>127:64</td><td>Zone Descriptor 0: Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).</td></tr> <tr> <td>191:128</td><td>Zone Descriptor 1: Contains the Zone Descriptor for the second zone reported, if any.</td></tr> <tr> <td>...</td><td>...</td></tr> <tr> <td>$((n+1)*64)+63:(n+1)*64$</td><td>Zone Descriptor n: Contains the Zone Descriptor for the last zone reported, if any.</td></tr> </tbody> </table>		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.
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.														
Source: <i>Id.</i> at 26.															

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
	Eh	Full	
Fh	Offline		
3:0	Reserved		

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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	

Source: *Id.* at 27–28.

Samsung's Accused Instrumentalities are able to write to the blocks (erase units) of a zone using a common page address. See:





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”



White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域 (Zone) 的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a), 当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b), 当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043]. Samsung’s Accused Instrumentalities are able to update the write pointer (information representing the extent to which the zone comprises host accessible pages available for writing) in response to writing to the set of blocks (erase units) of the 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. Source: ZNS Specification, at 9. 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.
--	--	--

		Source: <i>Id.</i> 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  Source: <i>Id.</i>
22	The storage device of claim 21, wherein the information about the zones comprises a state of the zone.	Samsung's Accused Instrumentalities comprise the storage device of claim 21. See claim 21. Samsung's Accused Instrumentalities comprise the storage device of claim 21 wherein the information about the zones comprises a state of the zone. Samsung's Accused Instrumentalities are able to retain and output a state of the zone in accordance with claim 21. 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. Source: ZNS Specification, at 24.	
--	--	---	--

Figure 35: Report Zones Data Structure <table border="1"> <thead> <tr> <th>Bytes</th><th>Description</th></tr> </thead> <tbody> <tr> <td>07:00</td><td> 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. </td></tr> <tr> <td>63:08</td><td>Reserved</td></tr> <tr> <td>127:64</td><td>Zone Descriptor 0: Contains the Zone Descriptor for the first zone reported, if any (refer to Figure 37).</td></tr> <tr> <td>191:128</td><td>Zone Descriptor 1: Contains the Zone Descriptor for the second zone reported, if any.</td></tr> <tr> <td>...</td><td>...</td></tr> <tr> <td>$((n+1)*64)+63:(n+1)*64$</td><td>Zone Descriptor n: Contains the Zone Descriptor for the last zone reported, if any.</td></tr> </tbody> </table>		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.
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.														
Source: <i>Id.</i> at 26.															

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
		Eh	Full
		Fh	Offline
		3:0	Reserved

Source: *Id.* at 27.


23 The storage device of claim 22, wherein the state of the zone is a full state when the zone has no host accessible

Samsung's Accused Instrumentalities comprise the storage device of claim 22. See claim 22.

Samsung's Accused Instrumentalities comprise the storage device of claim 22 wherein the state of the zone is a full state when the zone has no host accessible pages available for writing the zone.

See:

pages available for writing the zone.



Source: ZNS Specification, at 12.

Transition ZSIO:ZSF: The zone shall transition from the ZSIO:Implicitly Opened state to the ZSF:Full state:


- as a result of successful completion of a Zone Management Send command with a Zone Send Action of Finish Zone;
- as a result of a write operation that writes one or more logical blocks that causes the zone to reach its writeable zone capacity;
- due to a Zone Active Excursion (refer to section 5.6); and
- as a result of the zoned namespace becoming write protected (refer to the Namespace Write Protection section in the NVM Express Base Specification).

		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). Source: <i>Id.</i> at 13.
24	The storage device of claim 23, wherein the state of the zone is updated to the full state in response to the zone having no host accessible pages available for writing the zone after writing to the set of erase units of the zone of the zones using	Samsung's Accused Instrumentalities comprise the storage device of claim 23. See claim 23. Samsung's Accused Instrumentalities comprise the storage device of claim 23 wherein the state of the zone is updated to the full state in response to the zone having no host accessible pages available for writing the zone after writing to the set of blocks (erase units) of the zone of the zones using the common page address. See claims 21 and 23.

	the common page address.	
25	The storage device of claim 22, wherein the state of the zone is an empty state when all host accessible pages of the zone are available for writing.	Samsung's Accused Instrumentalities comprise the storage device of claim 22. See claim 22. Samsung's Accused Instrumentalities comprise the storage device of claim 22 wherein the state of the zone is an empty state when all host accessible pages of the zone are available for writing. See: Figure 7: Zone State Machine <pre> graph LR subgraph Active_Resources [Active Resources] ZSE((ZSE:Empty)) ZSF((ZSF:Full)) ZSC((ZSC:Closed)) ZSE <--> ZSF ZSE --> ZSC ZSC --> ZSE end subgraph Open_Resources [Open Resources] ZSIO((ZSIO: Implicitly Opened)) ZSEO((ZSEO: Explicitly Opened)) ZSIO --> ZSEO end ZSC <--> ZSIO ZSC <--> ZSEO ZSIO --> ZSRO((ZSRO:Read Only)) ZSRO --> ZSO((ZSO:Offline)) ZSEO --> ZSO </pre> The diagram illustrates the Zone State Machine. It features two main containers: 'Active Resources' and 'Open Resources'. 'Active Resources' contains three states: ZSE:Empty, ZSF:Full, and ZSC:Closed. ZSE:Empty and ZSF:Full are connected by a bidirectional arrow. ZSE:Empty has a solid arrow pointing to ZSC:Closed, and ZSC:Closed has a dashed arrow pointing back to ZSE:Empty. 'Open Resources' contains two states: ZSIO: Implicitly Opened and ZSEO: Explicitly Opened, connected by a downward arrow. ZSC:Closed has bidirectional arrows connecting it to both ZSIO: Implicitly Opened and ZSEO: Explicitly Opened. Outside these containers, ZSIO: Implicitly Opened has a solid arrow pointing to ZSRO:Read Only, which in turn points to ZSO:Offline. ZSEO: Explicitly Opened also has a solid arrow pointing to ZSO:Offline. Source: ZNS Specification at 12.

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



Source: *Id.* at 10.

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';

Source: *Id.* at 11.

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.

Source: *Id.* at 12.

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.

Source: *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.

Source: *Id.*

		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. Source: <i>Id.</i> at 14. All logical blocks in a zone shall be marked as deallocated when the zone is in the ZSE:Empty state. Source: <i>Id.</i> at 15.
27	The storage device of claim 21, wherein the information about the zones comprises logical addresses associated with respective ones of a plurality of zones of the zones, and information representing an extent to which each of the plurality of zones comprises host	Samsung's Accused Instrumentalities comprise the storage device of claim 21. See claim 21. Samsung's Accused Instrumentalities comprise the storage device of claim 21 wherein the information about the zones comprises logical addresses associated with respective ones of a plurality of zones of the zones, and information representing an extent to which each of the plurality of zones comprises host accessible pages available for writing. 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. Source: ZNS Specification, at 24.

accessible
pages
available for
writing.

15:08

Zone Receive Action Specific:			
Zone Receive Action	Description		
Report Zones	Reporting Options:		
	Value	Description	
	0h	List all zones.	
	1h	List the zones in the ZSE:Empty state.	
Extended Report Zones	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	
	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

Source: *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.

Source: *Id.* at 26.

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	
01	Bits	Description	
	7:4	Zone State (ZS): This field indicates the state of the zone.	
		Value	Definition
		0h	Reserved
		1h	Empty
		2h	Implicitly Opened
		3h	Explicitly Opened
		4h	Closed
		5h to Ch	Reserved
		Dh	Read Only
	Eh	Full	
Fh	Offline		
3:0	Reserved		

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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:												
	<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></table>	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.						
	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.												

		<table><tr><td></td><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><tr><td>07:04</td><td colspan="2">Reserved</td></tr><tr><td>15:08</td><td colspan="2">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 colspan="2">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 colspan="2">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 colspan="2">Reserved</td></tr></table>		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	
	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																			
		Source: <i>Id.</i> at 27–28.																		
28	The storage device of claim 21, wherein the zone is a minimum memory space associated with a management	Samsung’s Accused Instrumentalities comprise the storage device of claim 21. See claim 21. Samsung’s Accused Instrumentalities comprise the storage device of claim 21 wherein the zone is a minimum memory space associated with a management command from the host. Samsung’s Accused Instrumentalities are able to receive and execute the Zone Management Send command specifying a Reset Zone action (management command) which specifies at least a zone and is associated with a minimum memory space of at least a zone. See:																		

command
from the host.

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.																														
	<table><tr><th>Value</th><th>Description</th><th>Refer to section</th></tr><tr><td>00h</td><td>Reserved</td><td></td></tr><tr><td>01h</td><td>Close Zone: Close one or more zones.</td><td>3.4.3.1.1</td></tr><tr><td>02h</td><td>Finish Zone: Finish one or more zones.</td><td>3.4.3.1.2</td></tr><tr><td>03h</td><td>Open Zone: Open one or more zones.</td><td>3.4.3.1.3</td></tr><tr><td>04h</td><td>Reset Zone: Reset one or more zones.</td><td>3.4.3.1.4</td></tr><tr><td>05h</td><td>Offline Zone: Offline one or more zones.</td><td>3.4.3.1.5</td></tr><tr><td>06h to 0Fh</td><td>Reserved</td><td></td></tr><tr><td>10h</td><td>Set Zone Descriptor Extension: Attach Zone Descriptor Extension data to a zone.</td><td>3.4.3.1.6</td></tr><tr><td>11h to FFh</td><td>Reserved</td><td></td></tr></table>	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	
	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																														

Source: ZNS Specification, 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:

- 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:Implicitly 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. Source: <i>Id.</i> at 30–31.
29	The storage device of claim 21, wherein the logic is operable to cause the storage device to concurrently erase the set of	Samsung's Accused Instrumentalities comprise the storage device of claim 21. See claim 21. Samsung's Accused Instrumentalities comprise the logic to cause the storage device to concurrently erase the set of erase units of the zone. Samsung's Accused Instrumentalities are able to concurrently erase the set of blocks (erase units) of the zone which are on respective planes. See:

erase units of the zone.





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.

	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 <i>zone interleaving degree</i> 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.
--	--

Source: Han, at 149.

FIG. 7




		[0027] Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a nonvolatile memory device configured to retain data stored therein even when a power is turned off. For example, each of the plurality of storage devices 11a, 11b, 11c, and 11d may be (or include) a solid state drive. Though the electronic device 10 is illustrated as including four (4) storage devices, the plurality of storage devices is not limited thereto, and may include, e.g., more storage devices than illustrated. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may include a first storage area ZNS and a second storage area CNS. [0028] The first storage area ZNS may be used as a zoned namespace. Each of the plurality of storage devices 11a, 11b, 11c, and 11d may be configured to permit sequential writes (or only sequential writes), which are based on sequential logical addresses, in the first storage area ZNS. A portion of the first storage area ZNS, for example, a portion of a storage space may be used to implement Redundant Array of Inexpensive Disks (RAID). For example, the storage space may be used based on a RAID (e.g., RAID5 and/or RAID 4) standard.	
--	--	---	--

		[0113] FIG. 7 illustrates an example in which a first zone Z1 and a second zone Z2 are generated in the storage devices 11a, 11b, 11c, and 11d. Referring to FIGS. 1, 2, and 7, in at least one embodiment, one zone may correspond to four erase units. [0114] The storage devices 11a, 11b, 11c, and 11d may assign the first to fourth erase units EU1 to EU4 respectively included in the storage devices 11a, 11b, 11c, and 11d to the first zone Z1. Continuous logical addresses may be assigned to the first to 12th zone areas ZA1 to ZA12 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. [0115] Continuous logical addresses may be assigned to each of the first to fourth RAID areas RA1 to RA4 that are included in the first to fourth RAID stripes STRP1 to STRP4 defined in the first to fourth erase units EU1 to EU4. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
--	--	--	--

		[0116] A write pointer WP of the first storage device 11a may point out the first (or LSB) logical address among logical addresses where data are not written, for example, a logical address of the first zone area ZA1. Data may be written in a storage space of the first storage device 11a, which is pointed out by the write pointer WP. When the data is written in the storage space pointed out by the write pointer WP, the first storage device 11a may update the write pointer WP to point out the first (or least significant bit (LSB)) logical address belonging to the first storage device 11a from among the logical addresses of the first zone Z1. [0117] Likewise, the write pointer WP of the second storage device 11b may point out the second zone area ZA2, and the write pointer WP of the third storage device 11c may point out the third zone area ZA3. [0118] In each RAID stripe, when logical addresses of the RAID area are continuous to logical addresses of zone areas, the RAID area may also be pointed out by the write pointer WP. For example, when the logical addresses of the first RAID area RA1 are continuous to the logical addresses of the third zone area ZA3, the write pointer WP of the fourth storage device 11d may point out the first RAID area RA1. [0119] In each RAID stripe, when logical addresses of the RAID area are managed to be independent of (or are not continuous to) logical addresses of zone areas, the RAID area may not be pointed out by the write pointer WP. For example, the write pointer WP of the fourth storage device	
--	--	---	--

		11d may skip the first RAID area RA1 and may point out the sixth zone area ZA6 (or may jump to the sixth zone area ZA6). [0120] The storage devices 11a, 11b, 11c, and 11d may assign the fifth to eighth erase units EU5 to EU8 respectively included in the storage devices 11a, 11b, 11c, and 11d to the second zone Z2. Continuous logical addresses may be assigned to the 13th to 24th zone areas ZA13 to ZA24 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. [0121] Continuous logical addresses may be assigned to each of the fifth to eighth RAID areas RA5 to RA8 that are included in the fifth to eighth RAID stripes STRP5 to STRP8 defined in the fifth to eighth erase units EU5 to EU8. Logical addresses of one RAID area may be continuous or discontinuous to logical addresses of any other RAID area and may be continuous or discontinuous to logical addresses of a zone area.	
		Source: Samsung's '890 Patent Application at Fig. 7, [0027]–[0028], [0113]–[0121].	



Translation:

Left side of figure: “Figure 1”

White box under Zone 1 box: “Zone Mapping”

Top right caption:

- “Large Zone
- Spanning multiple dies
 - High performance”

		[0003] 当ZNS SSD的区域的尺寸较小时,区域跨越较少的晶粒(die),因此,针对区域的数据传输带宽较小;当ZNS SSD的区域的尺寸较大时,区域跨越较多的晶粒(die),因此,可以提高针对区域的数据传输带宽。当区域的尺寸较大时,该区域可能存储具有不同寿命的文件,因此该区域存储的文件可能无法被同时无效,从而导致该区域无法被及时重置回收,在这种情况下,ZNS SSD的存储空间利用率低。 <i>Translation:</i> “[0003] When the size of a ZNS SSD zone is small, the zone spans fewer dies, so the data transmission bandwidth for the zone is small; when the size of a ZNS SSD zone is large, the zone spans more dies, so the data transmission bandwidth for the zone can be improved. When the size of a zone is large, the zone may store files with different lifespans, so the files stored in the zone may not be invalidated at the same time, resulting in the zone being unable to be reset and recycled in time. In this case, the storage space utilization rate of the ZNS SSD is low.” [0042] 图1示出了针对ZNS SSD的区域(Zone)的数据传输带宽与区域的尺寸的关系的示意图。 [0043] 参照图1中的a),当ZNS SSD的区域的尺寸较大时,其可以跨越多个晶粒,这样能够保证较高的数据传输带宽。参照图1中的b),当区域的尺寸较小时,其仅跨越单个晶粒,这样针对区域的数据传输带宽小,从而降低数据传输性能。 <i>Translation:</i>
--	--	--

		“[0042] Fig. 1 shows a schematic diagram of the relationship between the data transmission bandwidth of a zone and the size of the zone for a ZNS SSD. [0043] Referring to a) in Fig. 1, when the size of a ZNS SSD zone is large, it can span multiple dies, which can ensure a higher data transmission bandwidth. Referring to b) in Fig. 1, when the size of a zone is small, it only spans a single die, which reduces the data transmission bandwidth for the zone and reduces the data transmission performance.” Source: Samsung’s ’816 Patent Application at Fig. 1a, [0003], [0042]–[0043].
30 <i>pr</i> <i>e</i>	The storage device of claim 21, wherein the storage device is operable to:	Samsung’s Accused Instrumentalities comprise the storage device of claim 21. See claim 21.
30 <i>a</i>	generate metadata indicating an amount of time since the zone was written to, and	Samsung’s Accused Instrumentalities comprise the storage device operable to generate metadata indicating an amount of time since the zone was written to. On information and belief, Samsung’s Accused Instrumentalities generate metadata including metadata indicating an amount of time since a zone was written to, including metadata such as age of data or data retention period. See:





FIG. 4

		[0067] The storage media 125 may have a retention period (which may depend on the usage history of the storage media 125, and, as such, may vary within the storage media 125); data that has been stored longer than the retention period (i.e., data having an age exceeding the retention period) may become unreliable and may be said to have expired. Data may be stored in the storage media 125 using an error correcting code, which may be a block code. When data is read from the storage media 125, a quantity of raw data, referred to as a code block, may be read from the storage media 125, and an attempt to decode it may be made. If the attempt fails, additional attempts (read retrials) may be made. With use, a portion, e.g., a block, of the storage media 125 may degrade to the point that the retention period becomes unacceptably short, and the block may be classified as a “bad block”. To avoid allowing this circumstance to render the entire storage media 125 inoperable, reserve space, referred to as “bad block management reserve space” may be present (e.g., included in each flash memory die or in each flash memory plane), and the controller 120, or another controller internal to the flash memory die or to the flash memory plane may begin to use a block in the reserve and cease to use the bad block.	
--	--	--	--

	[0089] FIG. 4 is a schematic data layout diagram of a storage device 110 in accordance with example embodiments of the disclosure. A block lifespan table 405 lists the retention period of each block in the storage device 110, and a mapping table 410 lists mappings from logical block address to physical block addresses. When the storage device 110 transitions to the sustainable read-only mode, the storage device 110 copies data from the faulty used space 415 (the physical space affected by the fault condition) to the rescue space 420. Data in the normal used space 425 is handled as it would be if the storage device 110 were in normal operating mode (instead of being in the sustainable read-only mode). Because the fault condition results in a loss of physical space, there is a corresponding reduction in the size of the used LBA space 430, with some of the previously available LBA space becoming unused LBA space 435. Firmware 440, executed by the controller 120, may cause the data to be moved, as described above, and the tables to be updated.	
--	---	--

[0090] FIG. 5A is a flow chart for a method of transitioning to the sustainable read-only mode, in accordance with example embodiments of the disclosure. At 505, the host 105 sends/receives data from storage device; at 510, the host 105 determines if an error has occurred in the storage device; at 520, the host 105 determines if an error happens with the storage device; at 525, the storage device 110 performs an internal diagnosis and transitions to the sustainable read-only mode; at 530, an FPGA may optionally monitor the nature of valid data on the storage device 110 (e.g., based on data type, or age); at 535, the storage device 110 relocates valid data to the rescue space when the age of the data approaches the retention period; and at 540, the storage device 110 may optionally transmit information about the status and contents of the rescue space to the host 105, which flushes the data. FIG. 5B is a flow chart showing details of

Source: Samsung's '781 Patent Application at Fig. 4, [0067], [0089]–[0090].



Source: 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	-

Source: *Id.* at 6.

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 implementation 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. Source: 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. Source: Samsung Application Note, at 2. 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. Source: Samsung Tech Blog.
--	--	---

		2.2.3. Reliability limitation A flash memory chip is unreliable in itself. Actually, during read and write operations, phenomena of bit inversion (<i>bitflips</i>) 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]. Source: Boukhobza, at 26.	
30b	in response to comparing the metadata to a threshold,	Samsung's Accused Instrumentalities comprise storage device operable to, in response to comparing the metadata to a threshold, provide an indication that the storage device recommends that the host send a second management command directed to the zone.	

provide an indication that the storage device recommends that the host send a second management command directed to the zone.	Samsung's Accused Instrumentalities are able to detect the need for management based on comparing the metadata to a threshold, provide a "Reset Zone Recommended" indication that the storage device recommends that the host send a Zone Management Send command specifying a Reset Zone action (second management command) directed to the 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 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'. Source: ZNS Specification, at 41–42.
--	---

Figure 37: Zone Descriptor Data Structure

Bytes	Description												
02	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>	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.
	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.												

Source: *Id.* at 28.

See also claim element [30a].