


NetApp INSIGHT 2023

Always On

数据存储与管理前沿技术系列课程


第二讲：ONTAP SAN特性

2023年11月24日



目录


Contents



- **ASA产品概览**
- **ONTAP SAN特性**

目录

Contents



- **ASA产品概览**

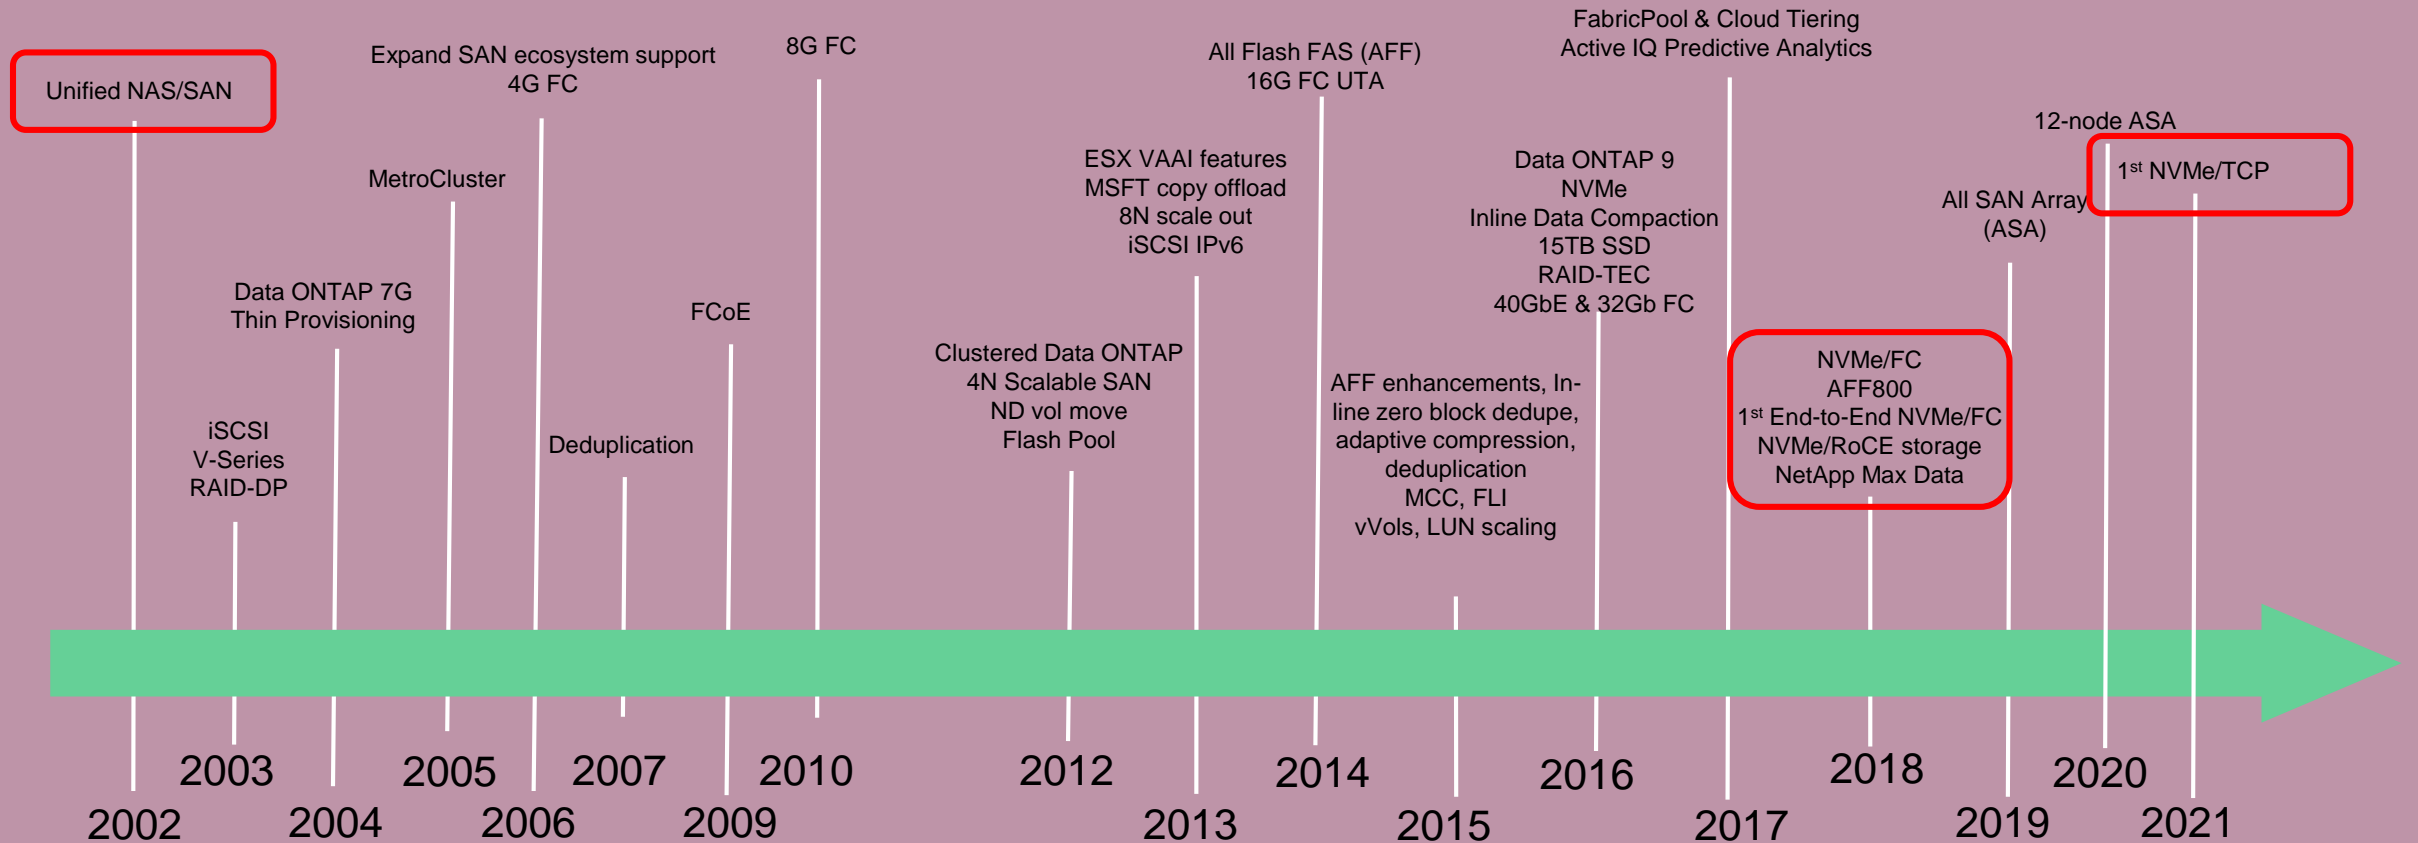
ONTAP SAN历程

Traditional Data Centers

Server Virtualization

Flash, SCM/PMEM

Cloud, Hybrid Cloud, Data Fabric




NetApp-公认SAN领域领导者

Technology innovation

1st End-to-end NVMe system	1st End-to-end 32Gb FC
1st SAN application acceleration with MAX Data	68% NetApp® AFF customers use SAN

Partner innovation

BROCADE
A Broadcom Limited Company
17-year partnership
with executive, engineering, field and technical integration



Market leadership

Fastest growing SAN vendor
Two years in a row

\$2B+
SAN revenues




Gartner identifies NetApp as leader in 2021 primary storage magic quadrant

“NetApp solid-state arrays have led in adoption of next-generation end-to-end NVMe capabilities and the host network connectivity, so that advanced customer infrastructures can fully leverage the performance of NVMe.”



NetApp® AFF All SAN Array (ASA)旨在满足关键任务应用程序需求，同时提供业界领先的云连接



从ONTAP®9.7开始简化SAN客户体验

Simple

- Block-only 专用及调优方案
- ONTAP System Manager简化

Speedy and bullet-proof

- 从主机对LUN进行对称A/A访问
- 无中断故障切换(到LUN的活动路径始终可用)

NetApp ASA package

- 在整个系列中提供高端SAN功能
 - ASA supported on all AFF controllers
 - All SAN Array set up from factory, point-of-sale solution
 - NVMe license可用，用于启用NVMe/FC和NVMe/TCP
- Software package licenses: ONTAP ONE, ONTAP BASE





ASA A 系列

用于任务关键型和性能要求高的数据库（Oracle、MS SQL Server、SAP HANA）和VMware基础架构

性能

提供卓越的性能，以改善客户体验，实现应用服务级别目标，并缩短实现结果的时间：

- 高带宽和 IOPS：集群中的 IOPS 高达数百万
- 一致的亚毫秒级延迟：低至 100 μ s



ASA C 系列

适用于对延迟要求不高的大容量关键业务应用，如数据库、VMware 以及备份和灾难恢复

效率、可持续性和更低的总体拥有成本

- 从HDD或混合闪存升级到高密度闪存
- 大幅减少存储占用空间和总体拥有成本
- 功耗最多可降低 70%，从而提高可持续发展能力

ASA A系列产品

广泛的性能和容量支持产品



HA系统

- | | | | | |
|--|--|---|---|---|
| <ul style="list-style-type: none"> • 15.2TB–547TB raw • Up to 2.1PB effective • 2U base w/24 drives | <ul style="list-style-type: none"> • 15.2TB–1.1PB raw • Up to 4.4PB effective • 2U base w/24 drives | <ul style="list-style-type: none"> • 15.2TB–14.6PB raw • Up to 58.6PB effective • 4U base + 2U24 drive shelf | <ul style="list-style-type: none"> • 22.8TB–6.6PB raw • Up to 26.3PB effective • 4U base w/48 drives | <ul style="list-style-type: none"> • 22.8TB–14.6PB raw • Up to 58.6PB effective • 8U base + 2U24 shelf |
|--|--|---|---|---|

集群

- | | | | | |
|--|--|--|---|---|
| <ul style="list-style-type: none"> • 91.2TB–3.3PB raw • Up to 13.1PB effective | <ul style="list-style-type: none"> • 91.2TB–6.6PB raw • Up to 26.4PB effective | <ul style="list-style-type: none"> • 91.2TB–87.6PB raw • Up to 351.6PB effective | <ul style="list-style-type: none"> • 136.8TB–39.6PB raw • Up to 157.8PB effective | <ul style="list-style-type: none"> • 138TB–87.6PB raw • Up to 351.2PB effective |
|--|--|--|---|---|

协议

NVMe/FC, NVMe/TCP, FCP, iSCSI

ASA C系列型号

广泛的性能和容量支持产品



ASA C250



ASA C400



ASA C800

HA 系统

- 122TBs – 734TBs 裸容量
- 有效容量高达 2.9PBs
- 2U 控制柜

- 122TBs – 1.46PBs 裸容量
- 有效容量高达 5.9PBs
- 4U 控制柜 + 2U 磁盘扩展柜

- 182TBs – 3.7PBs 裸容量
- 有效容量高达 14.7PBs
- 4U 控制柜

集群

- 122TBs – 4.4PBs 裸容量
- 有效容量高达 17.5PBs

- 122TBs – 8.8PBs 裸容量
- 有效容量高达 35.5PBs


- 182TBs – 22.2PBs 裸容量
- 有效容量高达 88PBs

协议

NVMe/FC, NVMe/TCP, FCP, iSCSI

目录

Contents



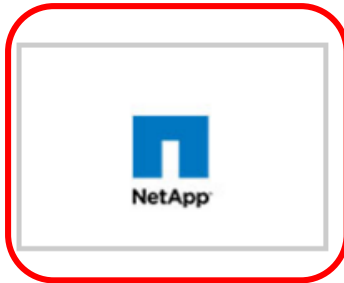
• ONTAP SAN特性

- ➔ □ NVMe
- ONTAP SAN特性

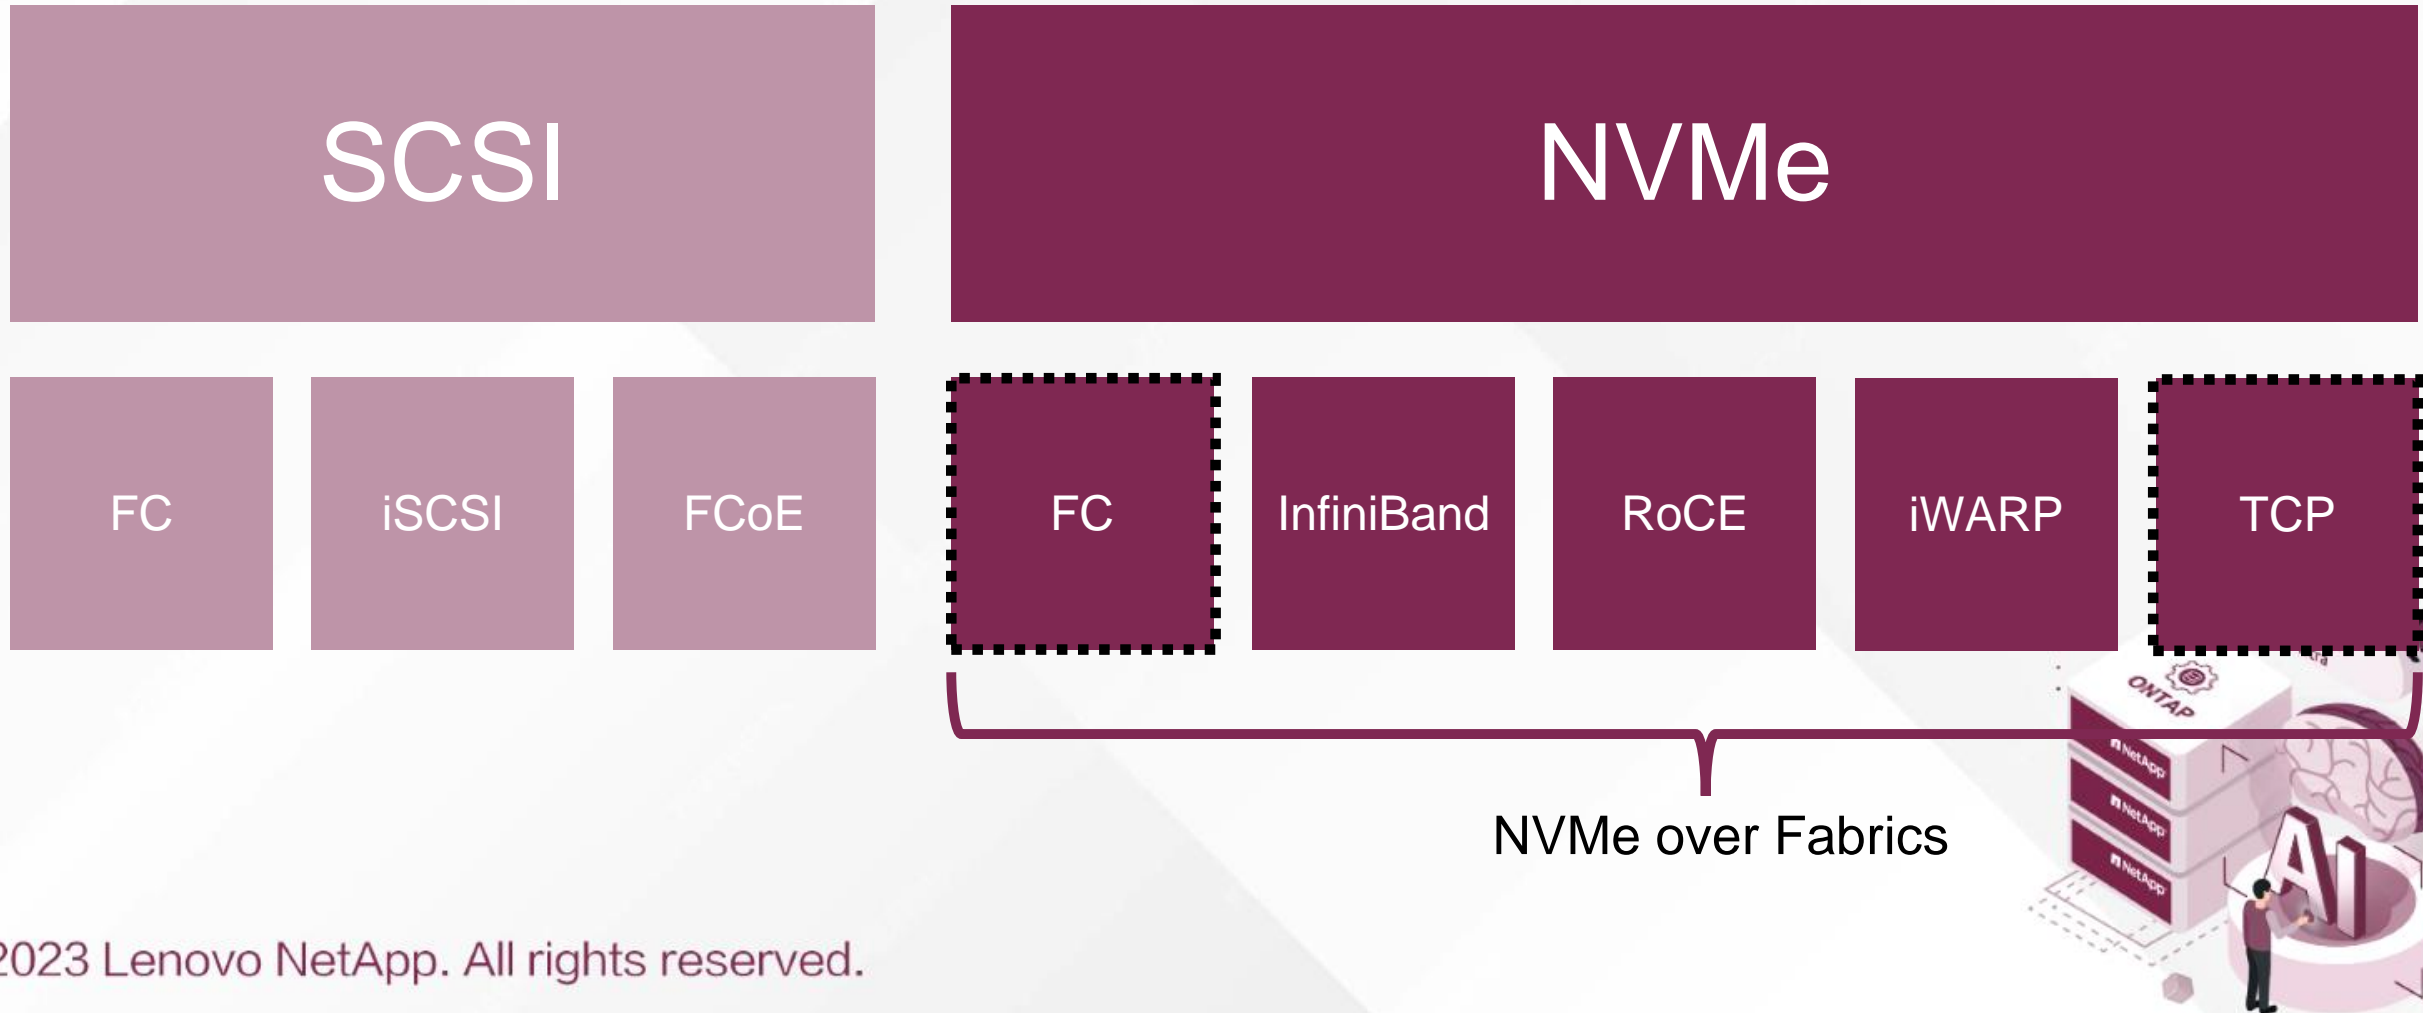
Promoter Information

[Home](#) > [About](#) > Promoter Information

PROMOTERS GROUP



什么是NVMe over Fabrics?



NVMe与SCSI命令集比较

NVMe IO Commands

1. Flush
2. Write
3. Read
4. Write Uncorrectable
5. Compare
6. Write Zeros
7. Dataset Management
8. Verify
9. Reservation Register
10. Reservation Report
11. Reservation Acquire
12. Reservation Release
13. Copy

Direct IO
commands:

NVMe	3
SCSI	35

SCSI IO Commands

- | | | | |
|----------------------------------|--|-----------------------------------|-----------------------------------|
| 1. TEST UNIT READY | 41. READ POSITION | 81. PERSISTENT RESERVE IN | 121. READ(12) |
| 2. REWIND | 42. SYNCHRONIZE CACHE(10) | 82. PERSISTENT RESERVE OUT | 122. SERVICE ACTION OUT(12) |
| 3. REQUEST SENSE | 43. LOCK UNLOCK CACHE(10) | 83. extended CDB | 123. WRITE(12) |
| 4. FORMAT | 44. READ DEFECT DATA(10) | 84. variable length CDB | 124. SERVICE ACTION IN(12) |
| 5. READ BLOCK LIMITS | 45. INITIALIZE ELEMENT STATUS WITH RANGE | 85. XDWRITE EXTENDED(16) | 125. ERASE(12) |
| 6. REASSIGN BLOCKS | 46. MEDIUM SCAN | 86. WRITE FILEMARKS(16) | 126. READ DVD STRUCTURE |
| 7. INITIALIZE ELEMENT STATUS | 47. COMPARE | 87. READ REVERSE(16) | 127. WRITE AND VERIFY(12) |
| 8. READ(6) | 48. COPY AND VERIFY | 88. Third-party Copy OUT commands | 128. VERIFY(12) |
| 9. WRITE(6) | 49. WRITE BUFFER | 89. Third-party Copy IN commands | 129. SEARCH DATA HIGH(12) |
| 10. SEEK(6) | 50. READ BUFFER | 90. ATA PASS-THROUGH(16) | 130. SEARCH DATA EQUAL(12) |
| 11. READ REVERSE(6) | 51. UPDATE BLOCK | 91. ACCESS CONTROL IN | 131. SEARCH DATA LOW(12) |
| 12. WRITE FILEMARKS(6) | 52. READ LONG(10) | 92. ACCESS CONTROL OUT | 132. SET LIMITS(12) |
| 13. SPACE(6) | 53. WRITE LONG(10) | 93. READ(16) | 133. READ ELEMENT STATUS ATTACHED |
| 14. INQUIRY | 54. CHANGE DEFINITION | 94. COMPARE AND WRITE | 134. SECURITY PROTOCOL OUT |
| 15. VERIFY(6) | 55. WRITE SAME(10) | 95. WRITE(16) | 135. SEND VOLUME TAG |
| 16. RECOVER BUFFERED DATA | 56. UNMAP | 96. ORWRITE | 136. READ DEFECT DATA(12) |
| 17. MODE SELECT(6) | 57. READ TOC/PMA/ATIP | 97. READ ATTRIBUTE | 137. READ ELEMENT STATUS |
| 18. RESERVE(6) | 58. REPORT DENSITY SUPPORT | 98. WRITE ATTRIBUTE | 138. READ CD MSF |
| 19. RELEASE(6) | 59. PLAY AUDIO(10) | 99. WRITE AND VERIFY(16) | 139. REDUNDANCY GROUP (IN) |
| 20. COPY | 60. GET CONFIGURATION | 100. VERIFY(16) | 140. REDUNDANCY GROUP (OUT) |
| 21. ERASE (6) | 61. PLAY AUDIO MSF | 101. PRE-FETCH(16) | 141. SPARE (IN) |
| 22. MODE SENSE (6) | 62. SANITIZE | 102. SYNCHRONIZE CACHE(16) | 142. SPARE (OUT) |
| 23. START STOP UNIT | 63. GET EVENT STATUS NOTIFICATION | 103. SPACE(16) | 143. VOLUME SET (IN) |
| 24. LOAD UNLOAD | 64. PAUSE/RESUME | 104. LOCK UNLOCK CACHE(16) | 144. VOLUME SET (OUT) |
| 25. RECEIVE DIAGNOSTIC RESULTS | 65. LOG SELECT | 105. LOCATE(16) | |
| 26. SEND DIAGNOSTIC | 66. LOG SENSE | 106. WRITE SAME(16) | |
| 27. PREVENT ALLOW MEDIUM REMOVAL | 67. XDWRITE(10) | 107. ERASE(16) | |
| 28. READ FORMAT CAPACITIES | 68. XPWRITE(10) | 108. SERVICE ACTION BIDIRECTIONAL | |
| 29. READ CAPACITY(10) | 69. READ DISC INFORMATION | 109. SERVICE ACTION IN(16) | |
| 30. READ(10) | 70. XDREAD(10) | 110. SERVICE ACTION OUT(16) | |
| 31. READ GENERATION | 71. XDWRITEREAD(10) | 111. REPORT LUNS | |
| 32. WRITE(10) | 72. SEND OPC INFORMATION | 112. ATA PASS-THROUGH(12) | |
| 33. SEEK(10) | 73. MODE SELECT(10) | 113. SECURITY PROTOCOL IN | |
| 34. LOCATE(10) | 74. RESERVE(10) | 114. MAINTENANCE IN | |
| 35. ERASE(10) | 75. RELEASE(10) | 115. MAINTENANCE OUT | |
| 36. READ UPDATED BLOCK | 76. REPAIR TRACK | 116. REPORT KEY | |
| 37. WRITE AND VERIFY(10) | 77. MODE SENSE(10) | 117. MOVE MEDIUM | |
| 38. VERIFY(10) | 78. CLOSE TRACK/SESSION | 118. PLAY AUDIO 12 | |
| 39. SET LIMITS(10) | 79. READ BUFFER CAPACITY | 119. EXCHANGE MEDIUM | |
| 40. PRE-FETCH(10) | 80. SEND CUE SHEET | 120. MOVE MEDIUM ATTACHED | |

NVMe与SCSI队列深度比较

NVMe

SCSI

- NVMe – 支持最大65,535队列，每个队列深度为65,535。
- SCSI – 单个队列深度为32/64/128
- 更多的队列意味着可以同时使用更多的处理器内核
- 更大的队列深度意味着可以用更少的空闲空间更有效地馈送网络管道



SATA vs. SAS vs. NVMe


INTERFACE	IOPS	THROUGHPUT	LATENCY	QUEUES	COMMANDS PER QUEUE
SATA	60,000 to 100,000	6 Gbps	below 1 millisecond (ms) to over 100 ms	1	32
SAS	200,000 to 400,000	12 Gbps	below 100 microseconds (μ s) to over 100 ms	1	256
NVMe	200,000 to 10,000,000	16 GBps (Gen3x16) 32 GBps (Gen4x16)	below 10 μ s to 225 μ s	65,535	64,000

Physical devices

Physical devices

目录


Contents



• ONTAP SAN特性

- NVMe优势
- ➔ □ ONTAP SAN特性

SAN高扩展-12-nodes



- ASA - HA对提供从主机到存储的对称双活路径。
- 所有节点作为单个集群进行管理。
- 所有无中断操作(NDO)、无中断升级(NDU) 支持。
- 跨不同型号， HA对之间数据流动。
- VOL/LUN移动和复制均支持。

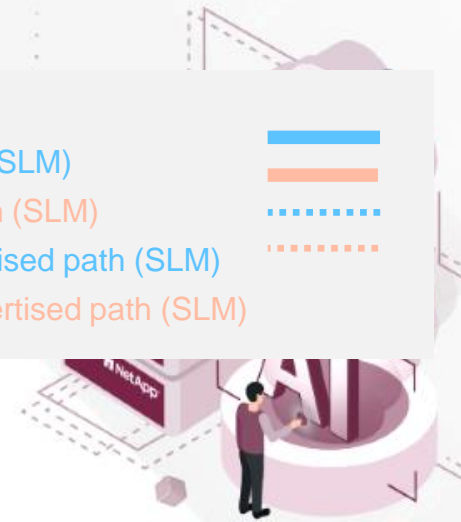
Legend

Blue Active Optimized (AO) advertised path (SLM)

Peach Active Optimized (AO) advertised path (SLM)

Blue Active Non-optimized (ANO) non-advertised path (SLM)

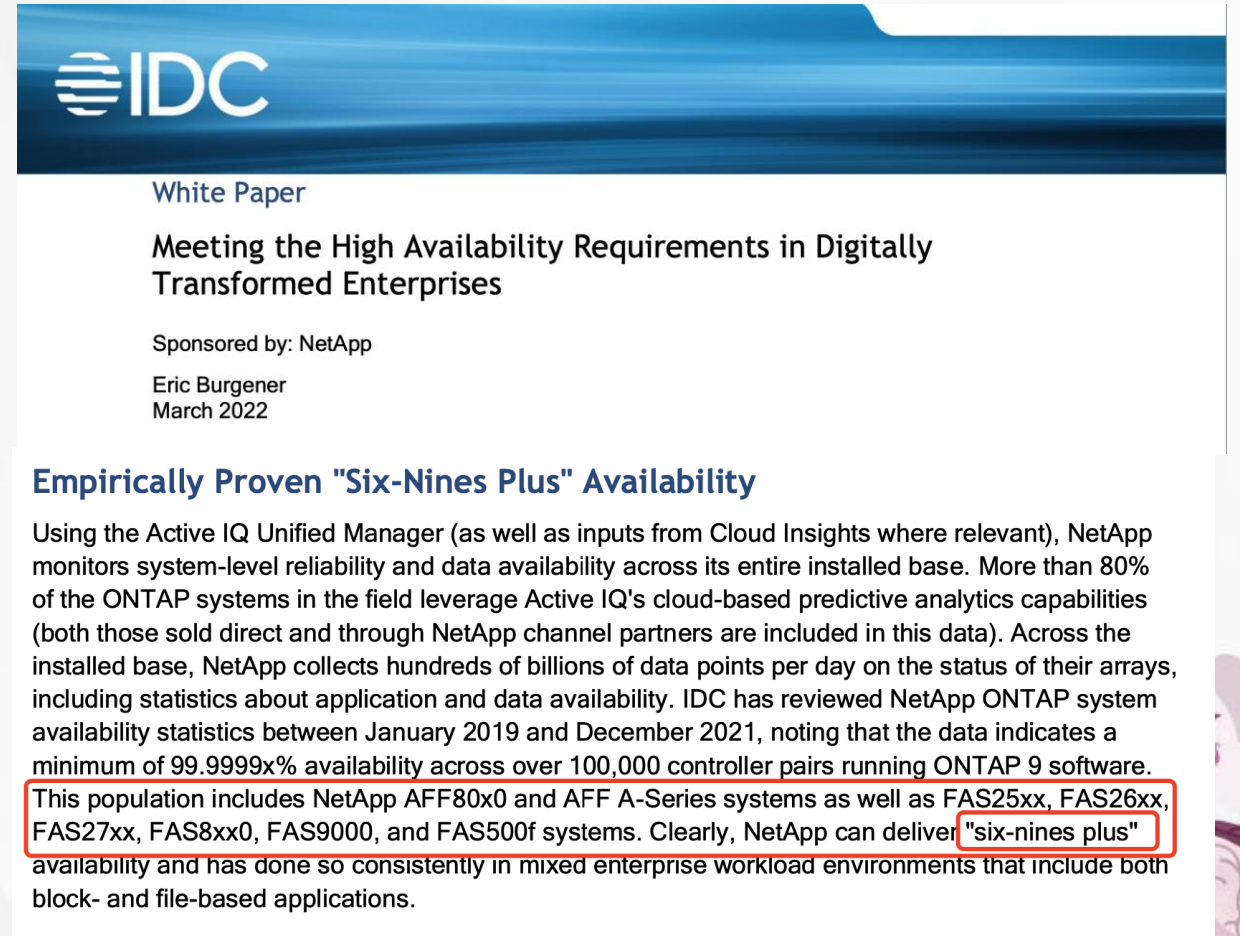
Peach Active Non-optimized (ANO) non-advertised path (SLM)



ASA高可靠-6个9

确保关键任务应用程序始终在线

- 对称双活路径到应用程序，实现高弹性的数据访问
- ASA存储旨在提供您运行业务所需的可用性—六个九(99.9999%)的正常运行时间
 - Fast controller failover
 - No maintenance windows
 - Nondisruptive SW and HW upgrades



IDC

White Paper

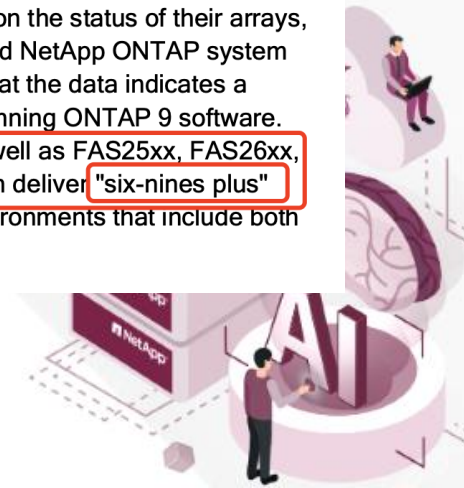
Meeting the High Availability Requirements in Digitally Transformed Enterprises

Sponsored by: NetApp

Eric Burgener
March 2022

Empirically Proven "Six-Nines Plus" Availability

Using the Active IQ Unified Manager (as well as inputs from Cloud Insights where relevant), NetApp monitors system-level reliability and data availability across its entire installed base. More than 80% of the ONTAP systems in the field leverage Active IQ's cloud-based predictive analytics capabilities (both those sold direct and through NetApp channel partners are included in this data). Across the installed base, NetApp collects hundreds of billions of data points per day on the status of their arrays, including statistics about application and data availability. IDC has reviewed NetApp ONTAP system availability statistics between January 2019 and December 2021, noting that the data indicates a minimum of 99.9999x% availability across over 100,000 controller pairs running ONTAP 9 software. This population includes NetApp AFF80x0 and AFF A-Series systems as well as FAS25xx, FAS26xx, FAS27xx, FAS8xx0, FAS9000, and FAS500f systems. Clearly, NetApp can deliver "six-nines plus" availability and has done so consistently in mixed enterprise workload environments that include both block- and file-based applications.




AFF 对比 ASA

AFF

通过两个控制器上的不等路径实现非对称 LUN 可用性


- 直接（所有者）路径经过主动优化 (AO)
- 间接（伙伴）路径为主动非优化路径 (ANO)



ASA

通过两个控制器上的相同路径实现对称 LUN 可用性

- 所有（所有者和伙伴）路径都是主动优化 (AO) 路径




加速关键工作负载

- 通过端到端NVMe存储加速企业应用程序
 - Extremely low latency (as low as 100 μ s)
 - Millions of IOPS in a cluster
- 通过在高性能和高密度存储上整合工作负载,降低数据中心成本



SAN功能增强-大LUN

- ONTAP® prior to 9.8




- Maximum LUN size: 16 TB

- Unified AFF-ONTAP 9.12.1P2 – larger LUNs



- ASA-ONTAP 9.8 – larger LUNs



LUN到namespace就地转换

Converting a LUN to a namespace

1. Display LUNs

```
tme-a700s-clus::> lun show
```

Vserver	Path	State	Mapped	Type	Size
svm0	/vol/testLUN/testLUN	online	mapped	linux	1GB

2. Unmap LUN to be converted

```
tme-a700s-clus::> lun unmap -vserver svm0 -path /vol/testLUN/testLUN -igroup new_15Mar21_tif5_igroup
```

3. Convert LUN

```
tme-a700s-clus::> vserver nvme namespace convert-from-lun -vserver svm0 -lun-path /vol/testLUN/testLUN
```

4. Map namespace to NVMe subsystem

```
vserver nvme subsystem map add -vserver svm0 -subsystem svm0_subsystem_909 -path /vol/testLUN/testLUN
```

5. Display new namespace

```
tme-a700s-clus::> vserver nvme namespace show
```

Vserver	Path	State	Size	Subsystem	NSID
Svm0	/vol/testLUN/testLUN	online	1GB	svm0_subsystem_909	00000001h



LUN到namespace就地转换

Converting a namespace to a LUN

1. Display new namespace

```
tme-a700s-clus::> vserver nvme namespace show
```

Vserver	Path	State	Size	Subsystem	NSID
Svm0	/vol/testLUN/testLUN	online	1GB	svm0_subsystem_909	00000001h

2. Unmap the namespace

```
vserver nvme subsystem map remove -vserver svm0 -subsystem svm0_subsystem_909 -path /vol/testLUN/testLUN
```

3. Convert namespace to a LUN

```
lun convert-from-namespace -vserver svm0 -namespace-path /vol/testLUN/testLUN
```

4. Map the LUN to an igroup

```
lun map -vserver svm0 -path /vol/testLUN/testLUN -igroup new_15Mar21_tif5_igroup -lun-id 20
```

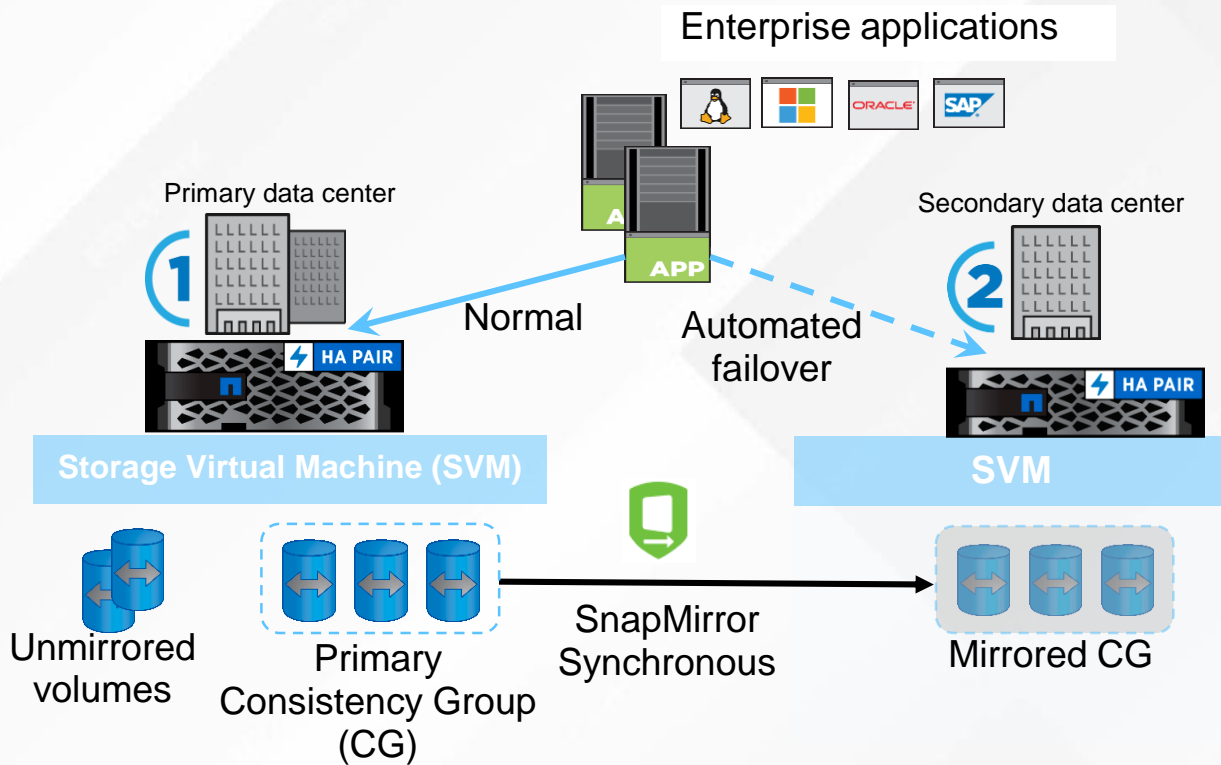
5. Display LUNs

```
tme-a700s-clus::> lun show -vserver svm0
```

Vserver	Path	State	Mapped	Type	Size
svm0	/vol/testLUN/testLUN	online	mapped	linux	1GB



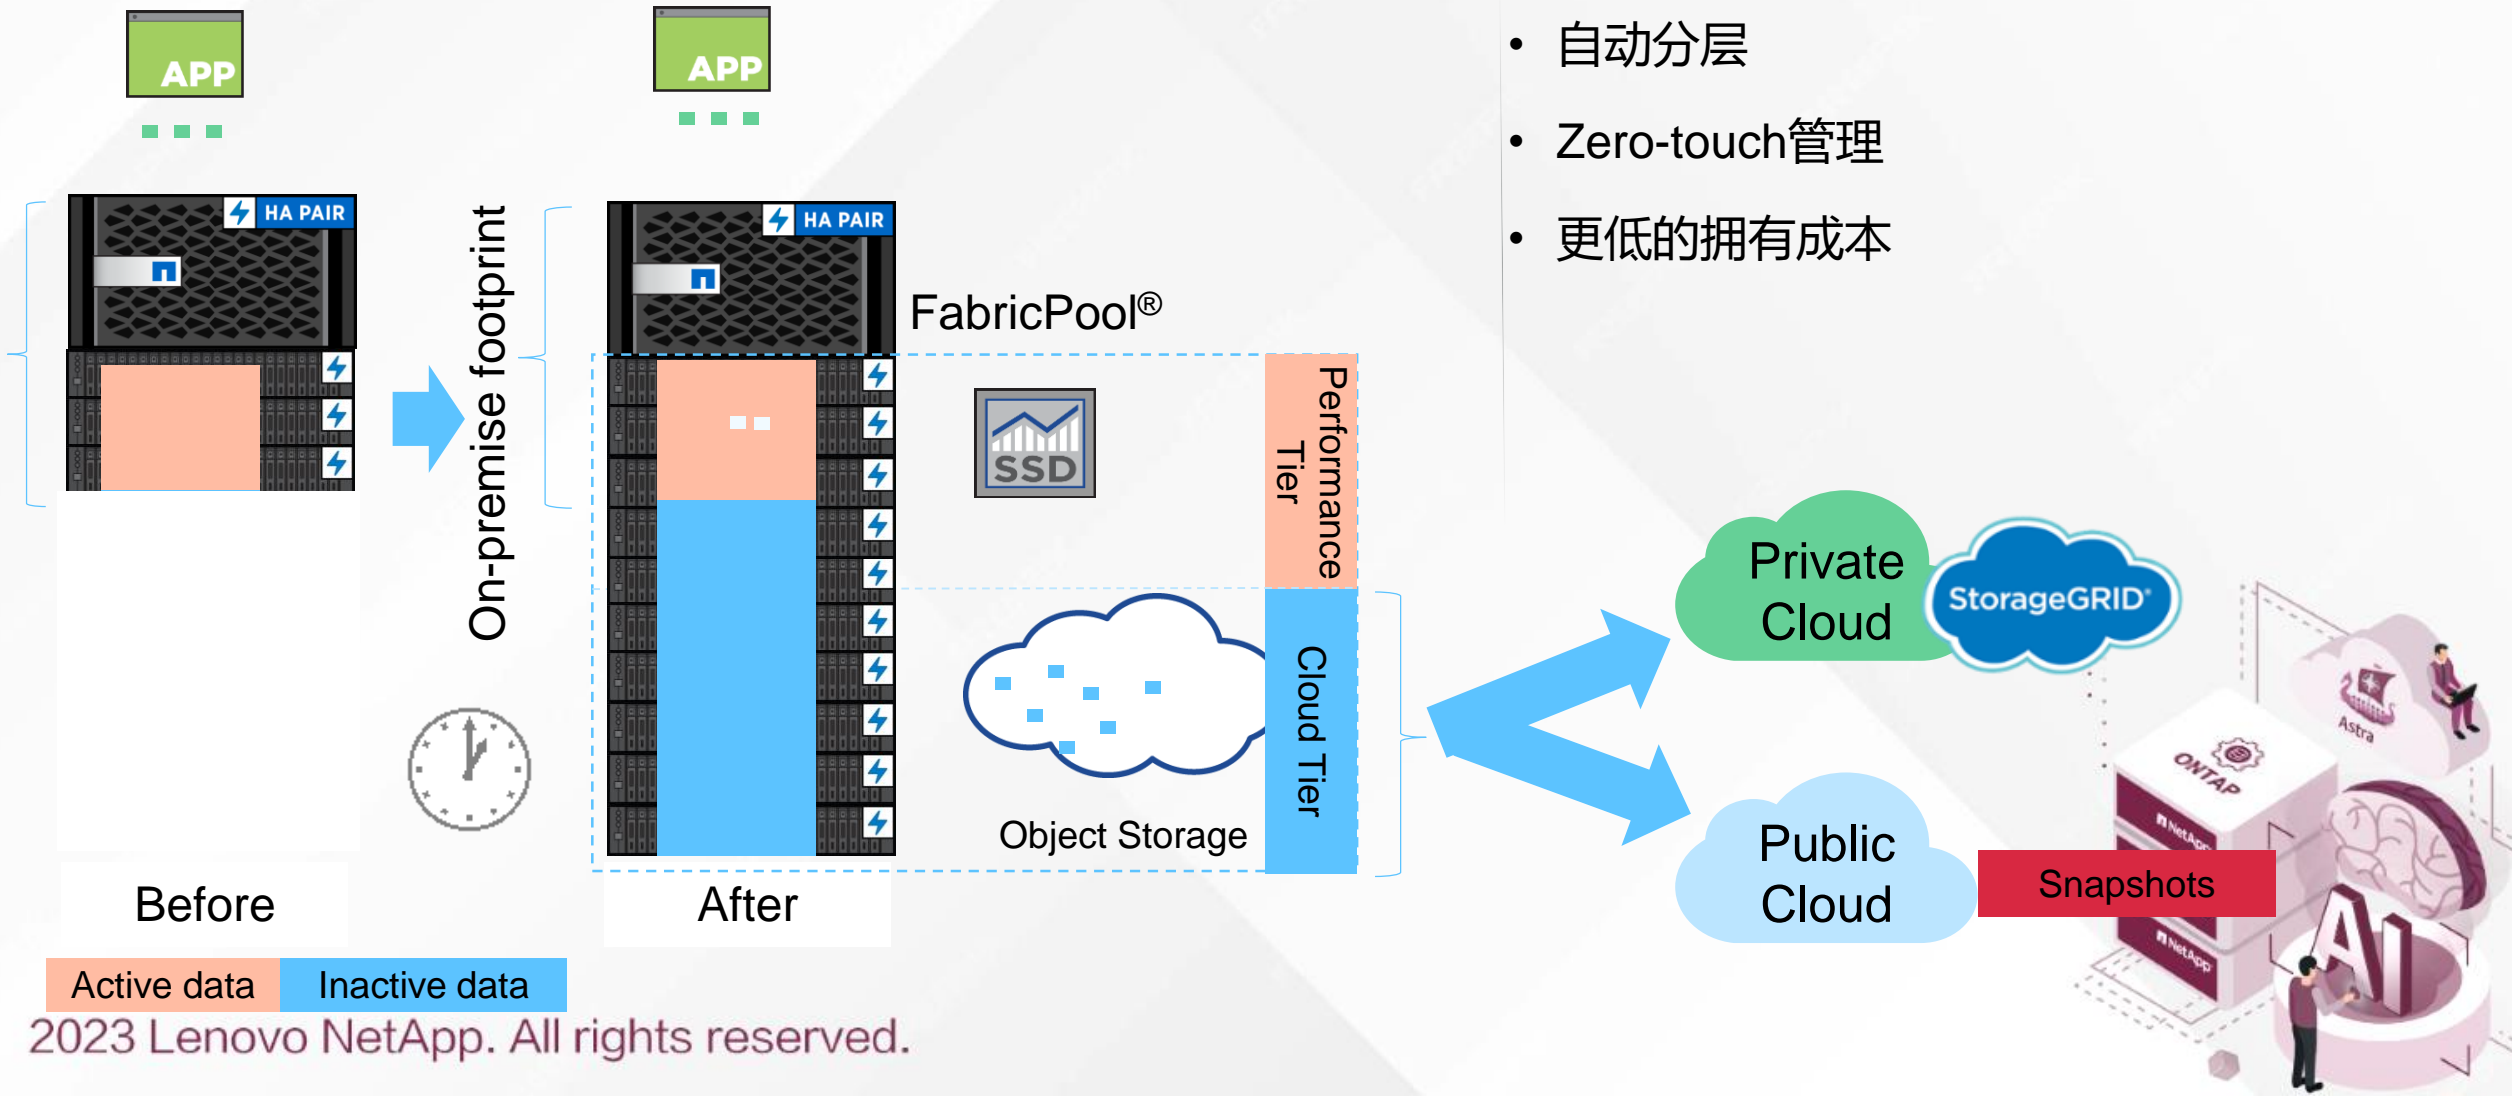
SAN数据保护-SMBC



- 满足SAN环境下的RTO=0与RPO要求
- 只能同类型部署
- 可支持仅保护关键应用
- 主备存储均可对外提供服务
- 主备存储档次可以不同
- 利用镜像副本进行开发和测试
- 10ms round-trip time (RTT) between sites
- 仅支持双节点集群



分层到云的FabricPool



满足各种需求的数据存储 - ONTAP

成本最低
次要用例

性价比最高

获得最佳性能
在一级工作负载上

混合闪存

闪存

性能闪存

统一

FAS



AFF C 系列



AFF A 系列



块优化

ASA C 系列

新



ASA A 系列



ONTAP

全面的数据管理软件，为文件、数据块和对象提供自动化、效率、数据保护和安全管理



谢谢!

联想凌拓官方联络方式

官方网站: www.lenovonetapp.com

服务热线: 400-828-3001 (呼叫中心)

400-116-0099 (销售热线)

官方社交平台账号:



联想凌拓
官方微信



Bilibili联想凌拓
空中沙龙



联想凌拓
渠道微信

