


# ASA 特性介绍与SAN架构评估

谭阳 | Service  
June 2024

# 目录

## Contents



- **ASA 特性介绍**
- **SAN 架构评估**

# ALL SAN Array

## All-flash 极致性能

### END-TO-END NVME

- NVMe SSDs
- NVMe-oF SUPPORT



### CONSISTENTLY HIGH IOPs



### SCALE UP AND SCALE OUT WITH CONSISTENT PERFORMANCE



## High availability 高可用

NetApp ASA 存储旨在提供您运营业务所需的可用性 - 六个九 (99.9999% 的正常运行时间)

## 业界先进的存储效率技术

使用 SAN 协议可以达到超过 4:1 存储效率\*

# License NetApp ONTAP One for SAN

一个许可证即可提供所有 ONTAP 功能

- 业内最全面的存储软件套件
- 成本可预测性和终极功能
- 支持所有ASA系统

## NetApp ONTAP One for SAN

Integrated remote backup/disaster recovery

Tamper-proof Snapshots

Application-consistent NetApp Snapshot™ copies

Multi-admin verification

Regulatory-compliant data retention

In-flight and data-at-rest encryption

Multifactor admin access

AI-informed predictive analytics and corrective action

Quality of service (QoS) workload control

Intuitive GUI, REST APIs, and automation integration

Automatic data tiering

Inline data compression, deduplication, and compaction

Symmetric active-active multipathing

FC, iSCSI, NVMe-oF, NVMe/TCP

## FC and iSCSI LIFs

- 每个SVM有自己的 IQN (iSCSI) or WWNN/WWPN (FC) and a service
  - 该服务必须正在运行, LIF 才能显示为已管理并可用
- FC LIFs
  - 每个端口都有自己的 WWPN, 并使用 NPIV 在物理 HBA 端口之上创建虚拟接口
  - 物理端口 WWNN/WWPN 看起来像 - 50:0a:09:80:80:61:d3:93
  - 逻辑接口 (LIF) WWNN/WWPN 看起来像 - 20:09:00:a0:98:9e:f8:55
  - NetApp 的组织唯一标识符 - 00:0A:98
  - 从集群节点的管理 MAC 地址派生的最后三对 LIF WWPN
- iSCSI LIFs
  - 每个都有自己的 IP 地址, target portal
  - Assigned to an IPspace
  - 只能配置 iSCSI, 不能配置其他协议 (SAN LIF 不会在接管时迁移\*, NAS LIF 会迁移)

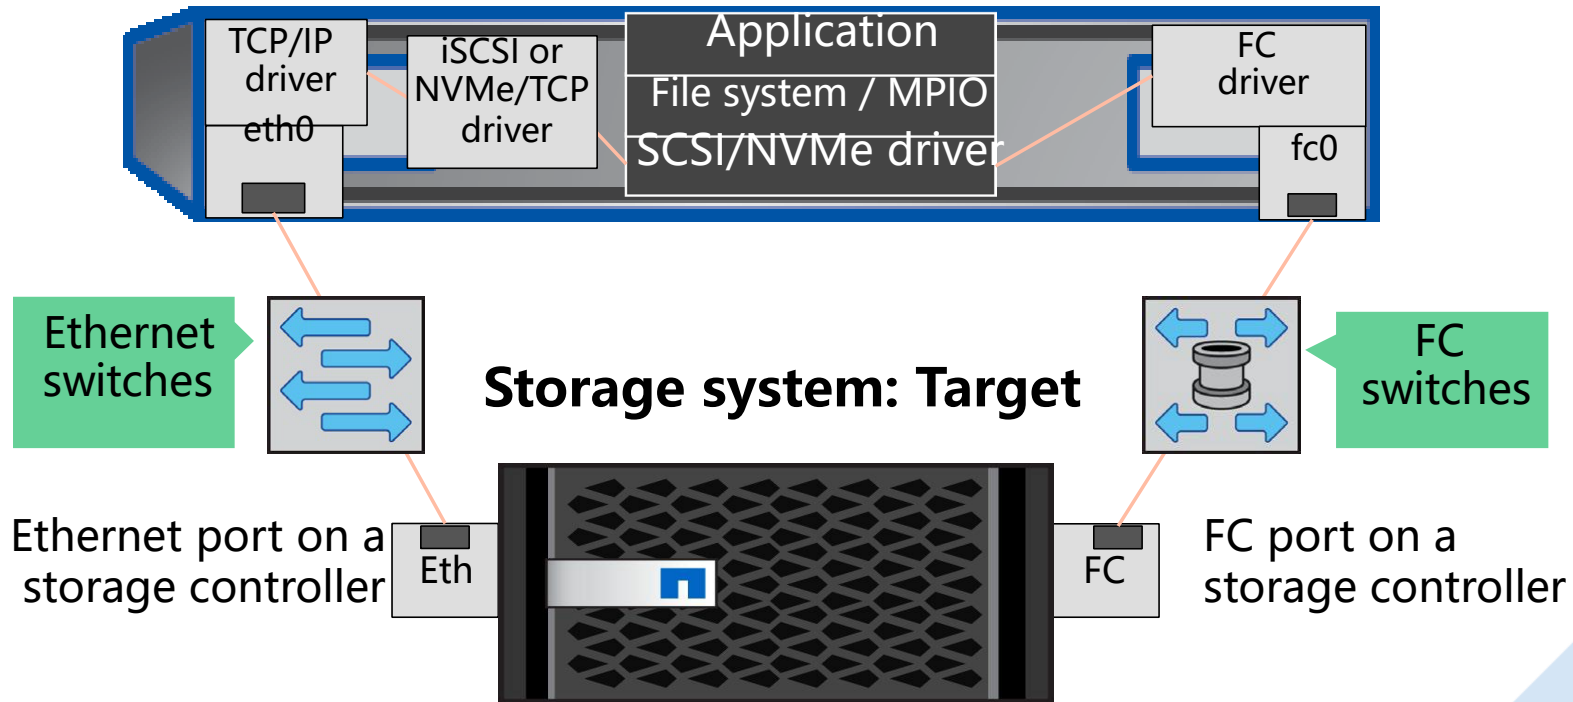
## NVMe LIFs

- 每个SVM都有自己的 NVMe Qualified Name (NQN) and a service
  - 该服务必须正在运行, LIF 才能显示为已管理并可用
- NVMe LIFs
  - 每个端口都有自己的 WWPN, 并使用 NPIV 在物理 HBA 端口之上创建虚拟接口
  - 物理端口 WWNN/WWPN 看起来像 - 50:0a:09:80:80:61:d3:93
  - 逻辑接口 (LIF) WWNN/WWPN 看起来像 - 20:09:00:a0:98:9e:f8:55
  - NetApp 的组织唯一标识符 - 00:0A:98
  - 从集群节点的管理 MAC 地址派生的最后三对 LIF WWPN
  - NVMe Namespace 必须映射到至少具有单个主机 NQN 的子系统, 以便识别的主机可以访问该命名空间。这类似于将 LUN 映射到 igroup 并将 WWPN 或 IQN 添加到 igroup 以允许主机访问。

# Connecting an initiator to a target

Fabric or network attached

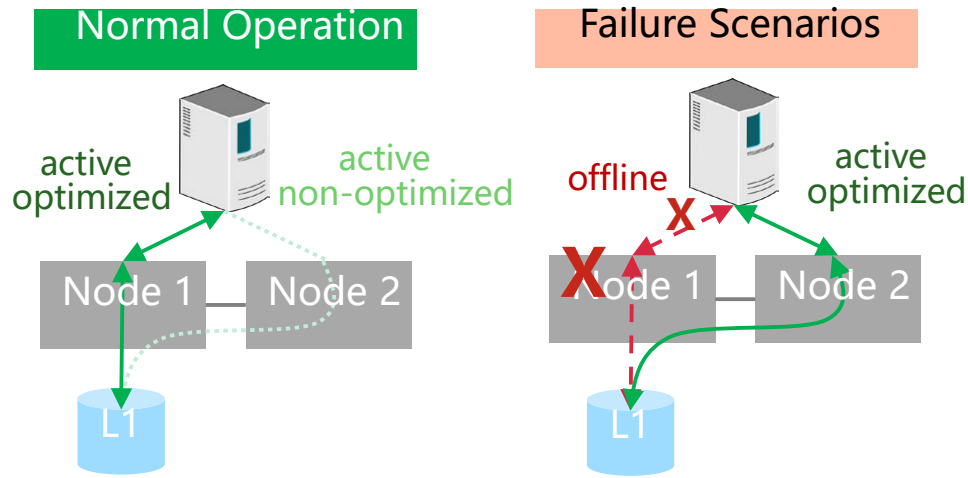
## Host: Initiator



# Asymmetric versus Symmetric access


## 故障转移期间的访问

### Unified AFF – asymmetric optimized/non-optimized



- 任何时候仅使用 1 个主机到 LUN 路径
- 当主控制器或优化路径发生故障时，数据流量会短暂暂停（通常为2-10秒）
- 这些暂停对于许多应用程序工作负载和客户来说是可以接受的，但并非全部

### AFF All SAN Array – symmetric active-active




- 使用所有可用的主机到 LUN 路径
- 当一条路径发生故障时，数据继续在其他路径上流动——数据持续可用
- 当存储控制器发生故障时，数据流量的暂停时间很短（通常 < 2秒）



# Asymmetric HA architecture

NetApp AFF unified array for business-critical enterprise applications



S1 active optimized paths:  
S1>FabA>N1    S1>FabB>N1

# Asymmetric HA host-side view

Unified steady state and in takeover (transitioning)

Unified steady state 2x AO, 2x ANO:

[root@client ~]# **multipath -ll**

```
3600a098038313546673f506f6b484464 dm-10
NETAPP,LUN C-Mode
size=20G features='3 queue_if_no_path
pg_init_retries 50' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|- 10:0:2:1 sdw 65:96 active ready running
|- 11:0:2:1 sdac 65:192 active ready running
-+- policy='service-time 0' prio=10 status=enabled
|- 10:0:3:1 sdz 65:144 active ready running
- 11:0:3:1 sdaf 65:240 active ready running
```


Unified takeover 2x active, 2x inactive paths:

```
3600a098038313546673f506f6b484464 dm-10
NETAPP,LUN C-Mode
size=20G features='3 queue_if_no_path
pg_init_retries 50' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|- 10:0:3:1 sdz 65:144 active ready running
|- 11:0:3:1 sdaf 65:240 active ready running
-+- policy='service-time 0' prio=0 status=enabled
|- 10:0:2:1 sdw 65:96 failed faulty running
- 11:0:2:1 sdac 65:192 failed faulty running
```

2 seconds of no I/O traffic,  
5 seconds of I/O degradation

# Symmetric HA architecture

NetApp AFF All SAN Array (ASA) for business-critical enterprise applications



optimized with only a small overhead

S1 active optimized paths:

S1 > FabA > N1      S1 > FabB > N1

S1 > FabA > N2      S1 > FabB > N2

# Symmetric HA host-side view

## ASA steady state and during takeover

ASA steady state:

```
[root@client ~]# multipath -ll
```

```
3600a0980383041336924506d5868735a dm-5
NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path
pg_init_retries 50' hwhandler='1 alua' wp=rw
+-+ policy='service-time 0' prio=50 status=active
| - 10:0:0:2 sdd 8:48 active ready running
| - 10:0:1:2 sdh 8:112 active ready running
| - 11:0:0:2 sdl 8:176 active ready running
| - 11:0:1:2 sdp 8:240 active ready running
```

ASA takeover 2x active, 2x inactive paths:

```
[root@client ~]# multipath -ll
```

```
3600a0980383041336924506d5868732f dm-4
NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path
pg_init_retries 50' hwhandler='1 alua' wp=rw
+-+ policy='service-time 0' prio=50 status=active
| - 10:0:0:1 sdc 8:32 failed faulty running
| - 10:0:1:1 sdg 8:96 active ready running
| - 11:0:0:1 sdk 8:160 active ready running
| - 11:0:1:1 sdo 8:224 failed faulty running
```

0 seconds of no I/O traffic,  
2 seconds of I/O degradation

# SAN 架构评估

# SAN 架构评估


如何确保SAN 交付策略与客户的业务目标保持一致?



确定 SAN 工作负载的当前状态、挑战、差距和期望的未来状态



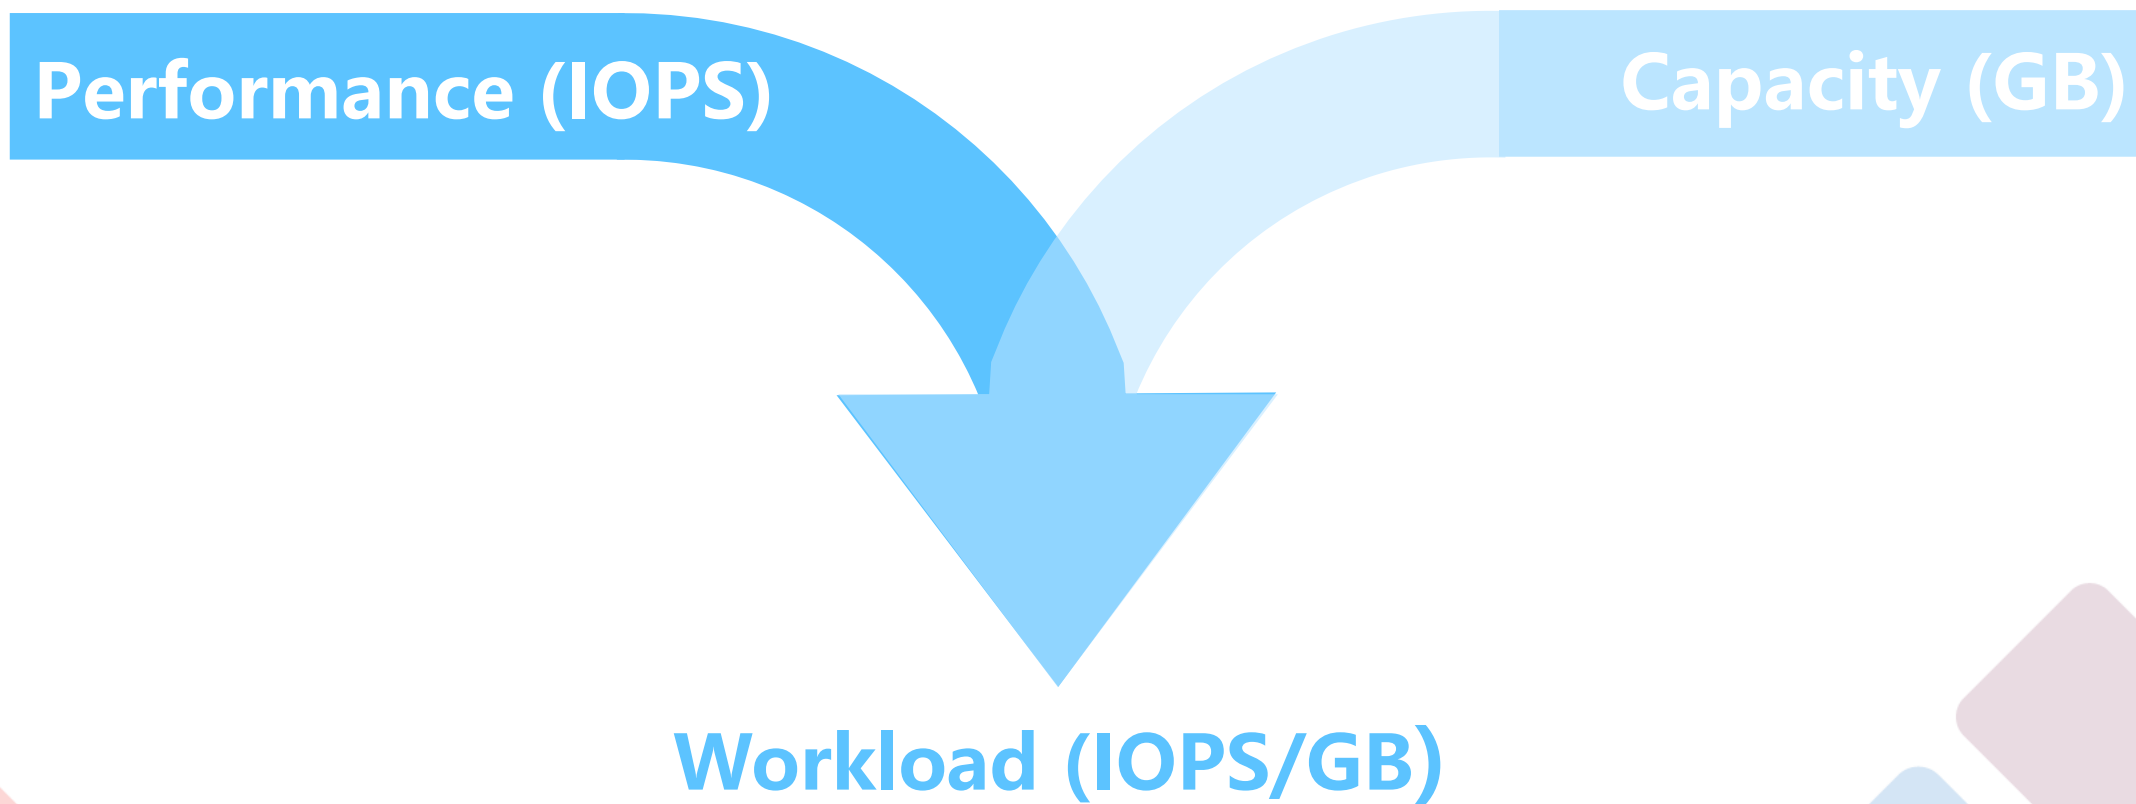
获取一组规范性建议，以加速实现客户的目标状态



创建明确的路径来继续客户的 SAN 之旅

# 存储设计挑战


平衡多种需求



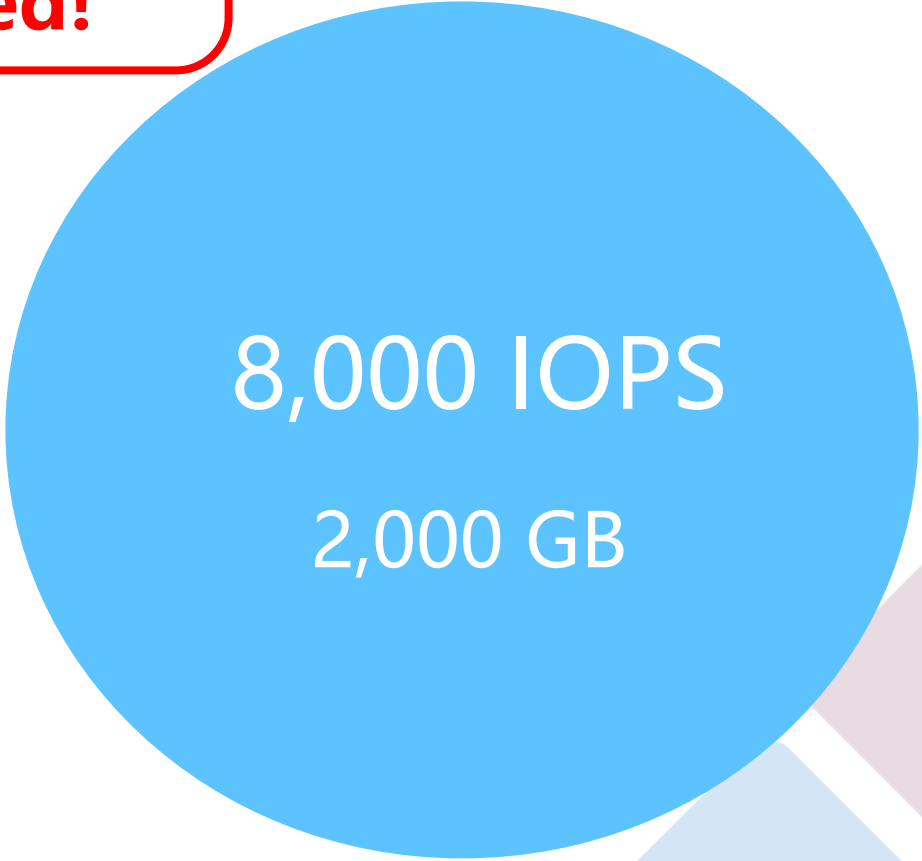
# IOPS versus IOPS/GB

Granularity

**4x Greater Performance Required!**



$$\frac{8,000 \text{ IOPS}}{500 \text{ GB}} = 16 \text{ IOPS/GB}$$



$$\frac{8,000 \text{ IOPS}}{2,000 \text{ GB}} = 4 \text{ IOPS/GB}$$



将最佳工具与自动化相结合

## Cloud Insights


- 基于SaaS
- 易于部署
- 专为云架构基础设施而设计
- 客户可以监控、优化和保护其所有基础设施

## I/O density report

- PS 架构师在最高级别的咨询服务 (SDW、CDW...等) 中使用了十年以上
- 准确测量和报告卷和 LUN 的容量与性能
- 提供有关卷/LUN 如何在未来状态解决方案 (本地、CVO...等) 上工作负载的见解
- 确定整合、分层和新产品购买的推荐

## 独特的知识产权和方法





- **I/O density analysis**
  - I/O 密度分析可描述数据速度 **Service-level design** 使交付符合要求
- **I/O budgeting**
  - I/O 预算计划容量的性能
- **Adaptive QoS policy management**
  - 自适应 QoS 策略管理可实现服务级别管理自动化、减少性能事件，并实现可预测的成本和性能
- **API-managed infrastructure**
  - API 管理的基础设施减少劳动力、支持 DevOps、提高混合云敏捷性



# Example - 服务级别协议和目标

基于行业验证的最佳实践

## Application-Aligned Storage Service Levels

	 Value	 Performance	 Extreme	 Ultra
Costs	\$	\$\$	\$\$\$	\$\$\$\$
Workload	Email, web, file shares, backup	Database and virtualized applications	Latency-sensitive applications	Top 2% of enterprise applications
Minimum SLA	0.125 IOPS/GB (128 IOPS/TB)	2 IOPS/GB (2,048 IOPS/TB)	6 IOPS/GB (6,144 IOPS/TB)	18 IOPS/GB (18,432 IOPS/TB)
Maximum SLO*	0.5 IOPS/GB (512 IOPS/TB)	4 IOPS/GB (4,094 IOPS/TB)	12 IOPS/GB (12,288 IOPS/TB)	36 IOPS/GB (36,864 IOPS/TB)
Latency Disk + Flash	17ms	2ms	1ms	1ms

1

2

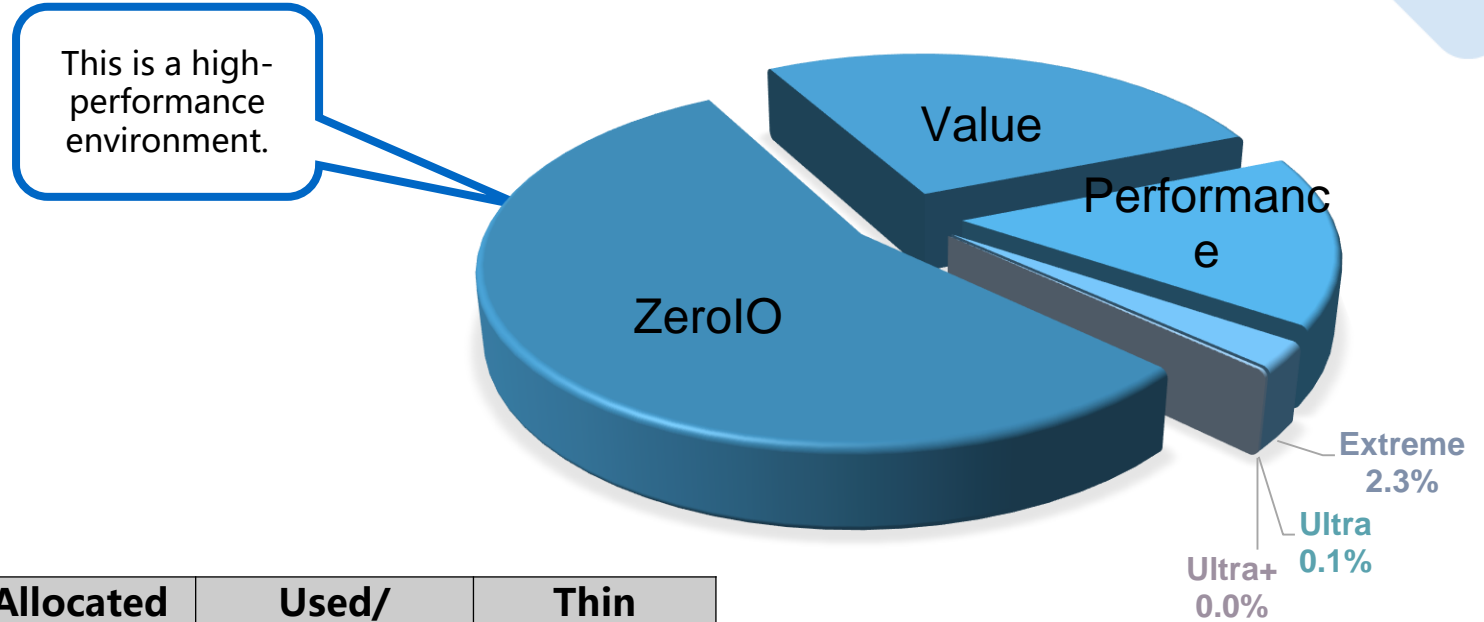
3

\*本表中的最大 SLO 源自 NetApp 服务部门多年来对客户工作负载的专有评估，并经过 NetApp 合作伙伴的验证

# "Volume Only" IO Density – Adaptive QoS "Volume Sort" View

All arrays

STANDARD SERVICE LEVEL RANGES  
(IO DENSITY - PEAK IOPS & USED CAPACITY)

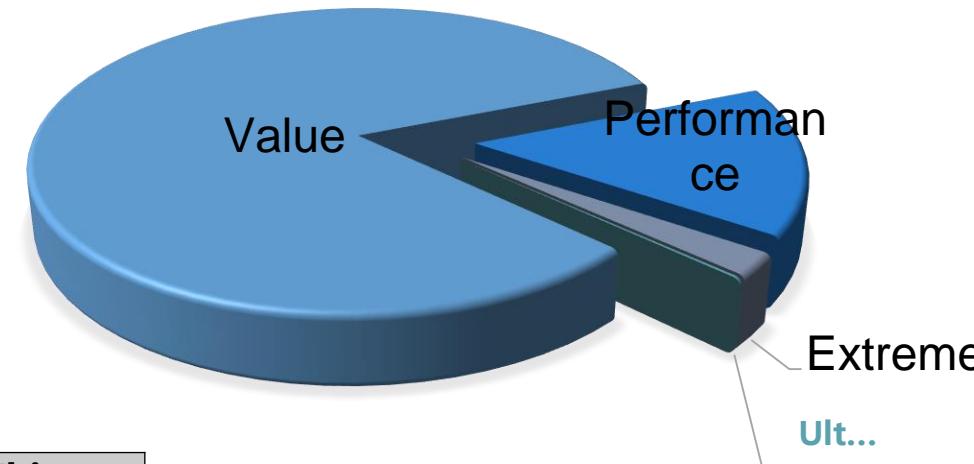


Standard Service Level	Service Level Ranges (IOPS/TB)	Volume Count	Used Capacity (TB)	Allocated Capacity (TB)	Used/Allocated Percentage	Thin Provision Gain (X:1)
ZeroIO	-	172	387.60	652.40	59.4%	1.68
Value	128-512	103	189.39	294.99	64.2%	1.56
Performance	2048-4096	79	109.61	196.99	55.6%	1.80
Extreme	6144-12288	32	16.36	36.70	44.6%	2.24
Ultra	18432-36864	8	1.01	2.14	47.0%	2.13
Ultra+	>36864	3	0.13	0.27	46.6%	2.15
<b>Grand Total</b>		<b>397</b>	<b>704.11</b>	<b>1,183.50</b>	<b>59.5%</b>	<b>1.68</b>

~82% of the Used percentage requires ≤512 IOPS/TB

## 基于活动的建议服务级别分布

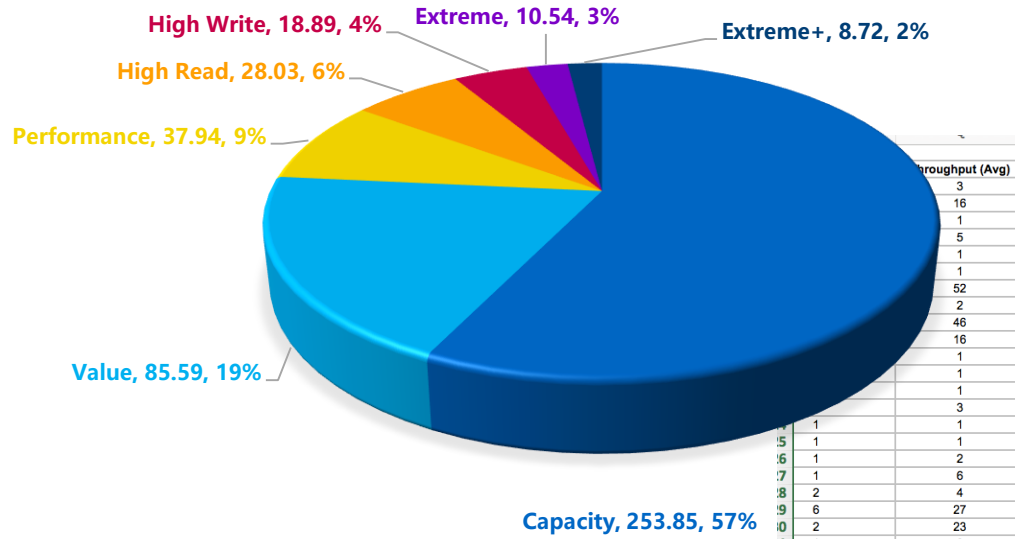
AQOS SERVICE LEVEL RANGES  
(IO DENSITY - PEAK IOPS & USED CAPACITY)



AQoS Service Level	Service Level Ranges (IOPS/TB)	Volume Count	Used Capacity (TB)	Allocated Capacity (TB)	Used/Allocated Percentage	Thin Provision Gain (X:1)
Value	128-512	275	577.00	947.39	60.9%	1.64
Performance	2048-4096	79	109.61	196.99	55.6%	1.80
Extreme	6144-12288	32	16.36	36.70	44.6%	2.24
Ultra	18432-36864	11	1.13	2.42	47.0%	2.13
<b>Grand Total</b>		<b>397</b>	<b>704.11</b>	<b>1,183.50</b>	<b>59.5%</b>	<b>1.68</b>


# VM IOdensity Analysis

## ALLOCATED VM WORKLOAD DEMAND



VM ID	Throughput (Avg)	Read IOPS (Avg)	Write IOPS (Avg)	IOPS (Avg)	IOPS (Peak)	I/O Density (Avg)	I/O Density (Peak)	Allocated (TB)	Used (TB)
3	222	21	243	2,528.66667	20,405	126,433	0.04	0.02	
16	1,567	24	1,590	2,695.66667	45,899	67,392	0.14	0.04	
1	16	1	16	2,900.8	3,791	96,693	0.03	0.03	
5	256	104	359	14,733.2	6,289	92,082	0.23	0.16	
1	75	71	145	2,412.4	4,832	34,463	0.07	0.07	
1	12	8	19	1,093.53333	5,510	54,677	0.02	0.02	
52	665	309	973	1,241.53333	26,641	31,038	0.04	0.04	
2	42	6	47	2,795.46667	3,325	27,955	0.1	0.1	
46	1,806	462	2,268	2,564.73333	19,719	21,373	0.12	0.12	
16	480	175	634	31,176.26667	2,556	39,970	0.78	0.78	
1	11	3	13	1,611.86667	2,609	80,593	0.02	0.02	
1	51	3	54	626.13333	17,651	20,871	0.12	0.03	
1	30	3	32	550.66667	3,368	27,533	0.02	0.02	
3	43	2	45	5,119.73333	2,022	39,383	0.39	0.13	
1	21	5	25	719	11,977	35,950	0.03	0.02	
15	1	10	4	1,505	2,367	75,250	0.02	0.02	
16	2	48	4	1,944.4	1,408	19,444	0.1	0.1	
7	6	495	35	530	1,161.53333	13,234	16,593	0.2	0.07
8	4	29	31	59	645.6	2,337	16,140	0.1	0.04
9	6	398	92	490	623.73333	14,527	15,593	0.05	0.04
10	2	23	28	378	3,531.06667	9,020	15,352	0.29	0.23
11	2	16	16	31	3,016.6	1,852	30,166	0.1	0.1
12	1	1,459	22	1,480	2,621.13333	11,633	14,562	0.2	0.18
13	3	29	30	58	1,049.93333	4,495	14,999	0.07	0.07

## VM count by vCPU Range



# SAN 架构咨询服务

## 数据采集




- 部署Cloud Insights软件\*
- 收集性能数据至少 14 至 30 天（最好是在月底或“热点”时期）

## 服务详情



- NAS/SAN/VMDK/虚拟机/云
- 已分配/已使用 GiB
- 读/写/平均/峰值 IOPS
- 读/写/平均/峰值延迟
- 读/写/平均/峰值吞吐量
- 重复数据删除和压缩增益
- 在输出工作簿中识别和过滤
- 非活动数据报告 (IDR)
- 服务水平数据和性能（例如标准、极限等）
- 充分利用 IO 密度报告和分析来确定工作负载性能特征

## 交付物



- 现有 SAN 性能初始分类的信息
- 基于从 Cloud Insights 收集的数据，使用 Fusion 优化 SAN 方案
- 明确 IT 运营转型的路线图和建议

\*Cloud Insights软件是基于SaaS平台，其数据采集的主机需要连接internet网络

# 谢谢!

智慧数据构建智能世界

官方网站: [www.lenovonetapp.com](http://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