A model of SAN performance

Dr. Alexandre Brandwajn PALLAS International Corporation San Jose, CA

Plan

- Motivation
- ◆ SAN and model
 - logical I/O
 - Storage Domain Server
 - PCI bus model
- Results example



Model goals

- Expected I/O performance for clients
- SAN dimensioning
 - number of servers
 - server power
 - server configuration
 - internal bottlenecks
- Operating range

SAN configuration

General configuration	SAN Dimensions and Fabric
<u>B</u> osts and Logical Devices	Configuration Fibre Channel
<u>E</u> xecute	
SUSI	Storage Domain Servers: 🚺 🚊
	Number of bosts: 1
	Total logical devices: 32
	OK Cancel Apply Help

Client model

- I/O workload
 logical devices
- "typical" workloads
- host computer
 - speed
 - SAN interface
 - future model

ost Ge	eneral Logical Device characteristics Host Bus Adapter characteristics
	Host number: 0
	Number of Logical Devices: 32 View Workload type
	CPU speed (MHz): 500 Custom
	Number of Host Bus Adapters: 2 - Skewed Friendly Unfriendly
	PCI bus
	Speed (MHz): 66 Width (bytes): 4
	Arbitration latency (ns): 10

I/O workload model

- Extensive Lead Time experience
 - I/O burstiness
 - transfer and seek variability
 - limited queue
 - read hit behavior
- Mapping
 - mirroring

I/O workload model

t General Logical Device characteristic:	8 Host Bus Adapter characteristics
Host: Host 0 Logical device	e number: 0 📑 Name: 🛛 🛨
🗖 All logical	devices
Workload	
☑ I/Os per second: 20	I/O rate variability: 1
Direct length (Kbytes): 16	Transfer variability:
I Fraction of zero seeks: 0.5	Seek time variability:
I Fraction of reads: 0.75	Hits on reads fraction: 0.85
☑ Queueing limit: 0	Buffer hit probability: 0.5
Configuration	
SDS number: 0	🏹 Physical device: 🛛 📑
🔽 Mirroring type	e: None 💌
Index of mirror SDS:	Mirror device:

Storage Domain Servers

- Set of physical devices
- Server CPUs
 - polling
- PCI buses
- Host Bus Adapters
 - PCI bus <==> Fibre Channel



I/O operations

• Read hit

- served from SDS cache
- Read miss
 - I/O Request Packet ==> Windows I/O Manager
 - data first copied into SDS cache
- ♦ Write
 - synchronous and asynchronous mirroring
 - mirror in another SDS
 - asynchronous write back

I/O operations



SDS model



- G/G/1/N approximation
- queueing and service time
- CPU and polling
 - overheads and threads
- PCI buses and HBAs
 - effective bandwidth
 - future: main bus, memory access



Storage Domain S	erver		×
More C	Iverheads	Physical Disks	1
PCI bus char	acteristics Host	Bus Adapter characteristics	
Spo deneral	CFU Uverneads & Polling	More Polling and Overneads	
	SDS number: 📔 🚊	1	
Nu	imber of CPUs:	4 .	
Nu	imber of PCI buses:	3	
Nu	umber of Host Bus Adapters:	9	
Nu	umber of physical disks:	16 🕂	
Nu	imber of device types:	1 🕂	
CF	PU speed (MHz):	550	

|--|

PCI bus characteristics	Host Bus Adapter characteristics	
More Overheads	Physical Disks	- j
SDS General CPU Overhead	s & Polling More Polling and Overhead	۶Ì
SDS number:	Reference speed (MHz): 550	
Average CPU times (ms) for		i.
Unproductive poll: 0.00266	Bead hit: 0.073	
Read miss: 0.073	Write direct: 0.087	
Write back:	Mirror primary:	
Mirror secondary:		
Productive poll t	hreshhold: 0.5	
Οκ	Cancel Apple Help	

SDS model - polling

More Overheads Physical Disks SDS General CPU Overheads & Polling More Polling and Overheads & Polling SDS number: Image: Comparison of the polling of the polling and Overheads & Polling More Polling and Overheads Overheads & Polling Numbers of polls for Image: Comparison of the polling of the polling of the polling Image: Comparison of the polling of	PCI bus ch	haracteristics	Host E	Bus Adapter chara	icteristics
SDS General CPU Overheads & Polling More Polling and Overhe SDS number: Image: SDS number:	More	e Overheads		Physical Dis	sks
SDS number: Image: SDS number: Numbers of polls for Read miss: Read hit: Image: Read miss: Image: Secondary: Write back: Image: Mirror primary: Image: Secondary: Fibre Channel Interface card overheads (ms) for Image: Secondary: Image: Secondary: Read hit: Image: Omega: Secondary: Image: Secondary: Image: Secondary: Write back: Image: Mirror primary: Image: Secondary: Image: Secondary:	DS General	CPU Overhead	ls & Polling	More Polling an	id Overhead
Numbers of polls for 5 Write direct: 3 Read hit: 2 Read miss: 5 Write direct: 3 Write back: 3 Mirror primary: 3 Secondary: 3 Fibre Channel Interface card overheads (ms) for 0.094 Write direct: 0.11 Write back: 0 Mirror primary: 0 Secondary: 0		SDS r	iumber: 🔟		
Read hit: 2 Read miss: 5 Write direct: 3 Write back: 3 Mirror primary: 3 Secondary: 3 Fibre Channel Interface card overheads (ms) for 7 Read hit: 0.094 Read miss: 0.094 Write direct: 0.11 Write back: 0 Mirror primary: 0 Secondary: 0	Numbers of p	olls for			
Write back: 3 Mirror primary: 3 Secondary: 3 Fibre Channel Interface card overheads (ms) for Read hit: 0.094 Read miss: 0.094 Write direct: 0.11 Write back: 0 Mirror primary: 0 Secondary: 0	Read hit:	2 Read m	iss: 5	- Write direct:	3
Fibre Channel Interface card overheads (ms) for Read hit: 0.094 Read miss: 0.094 Write direct: 0.11 Write back: 0 Mirror primary: 0 Secondary: 0	Write back:	3 Mirror pr	rimary: 3	– Secondary:	3
Read hit: 0.094 Read miss: 0.094 Write direct: 0.11 Write back: 0 Mirror primary: 0 Secondary: 0	Fibre Channe	Interface card ov	arheads (ms) f	or	
Write back: 0 Mirror primary: 0 Secondary: 0	Read hit:	0.094 Read mi	ss: 0.094	- Write direct:	0.11
	Write back:	0 Mirror pri	mary: 0	- Secondary:	0
PCI bus access: 0.004		PCI bus	access: 0.0	004	

SDS model - Host Bus Adapters

SDS General	CPU Overheads & Po	olling 📔 More Polli	ng and Overhea
More	Overheads	Physic	al Disks
PCI bus cha	racteristics	Host Bus Adapter	characteristics
SDS number	0	HBA number: 0	<u>.</u>
Model:	A2200f/66 🛛	PCI bus nun	iber: 0 🚊
Simplex band	lwidth (MB/s): 100	Width (by	tes) 8
Operation on I	PCI bus (latency in bus o	cycles)	
Acquisition	latency: 2	Initial Target Laten	су: 8
Subsequer	it transfers: 1	End of transfer:	1
	Burst size (data pl	hases): 32	
Ē	or 1 c		1 1.44

SDS model - PCI bus

SDS General	CPU Overhea	ids & Polling	More Polling an	d Overheads
More	Overheads	9 10 m	Physical Dis	:ks
PCI bus ch	aracteristics	Host	Bus Adapter chara	cteristics
SDS numbe	r: 0	PCI bus i	ndex: 0 📑	
Speed (M	Hz): 📴	Width (bytes): 8]
Arbitration	latency (ns):	0	HBAs: 3	J
Γ	OK	Cancel	Apply	Help

SDS model - physical disks

PCI bus characteristics Host Bus Adapter characteristics More Dverheads Physical Disk SDS number: Image: Characteristics Device characteristics Device number: Image: Characteristics Device type: Seagate Cheetah Avge read seek time (ms): 5.7	s
More Dverheads Physical Disk SDS number: Device number: Device number: Device characteristics Device type: Seagate Cheetah Revolution (ms): Avge read seek time (ms): 5.7 Avge write seek (ms):	5 98
SDS number: Device number: Device number: Device characteristics Device type: Seagate Cheetah Revolution (ms): Avge read seek time (ms): 5.7 Avge write seek (ms):	5.98
Device type: Seagate Cheetah 💌 Revolution (ms): Avge read seek time (ms): 5.7 Avge write seek (ms): 1	5.98
Avge read seek time (ms): 5.7 Avge write seek (ms):	0.00
Avge read seek time (ms): Avge write seek (ms):	CE
	6.0
Min data rate (MB/s); 18.3 Max data rate (MB/s);	28
······································	
Drive range	
10 to 15 🚍	

PCI bus model

- PCI bus model
 - effective data rate, wait for bus
- Host Bus Adapter
 - traffic fraction, wait for HBA
- fixed-point iteration between models



PCI bus model - adapter traffic







I/O Service and Queueing

Results - PCI bus



Conclusions

- Model of in-band SAN
- Logical I/O for clients
- Storage Domain Server model
 - CPU overheads and polling
 - PCI buses and Host Bus Adapters
 - physical devices
- ◆ Future: leverage Lead Time
 - controllers attached to SDS