

Dimensionando Hardware e Storage para SQL Server



Marcelo Sincic
Technical Consultant – ProConsult Brazil
Dell | IT Computing Services
MVP-MCT-MCITP-MCPD-MCTS
MCDBA-MCSA-MCAD
Sun Java Instructor
IBM CLP Notes
<http://marcelosincic.com.br>
msincic@gmail.com



#127 | RIO DE JANEIRO 2012

Agenda

- O que é IOPS
- Storage
- Network
- Processador

Storage

- Como dimensionar o storage para performance, tipos e custo

O que é IOPS

- Input Output Process per Second
Indica o numero de operações que um disco pode executar por segundo
- Calcular corretamente o IOPS é essencial em um ambiente onde se deseja performance
- O total de IOPS de um storage é medido pela soma dos IOPS dos discos individuais de um storage

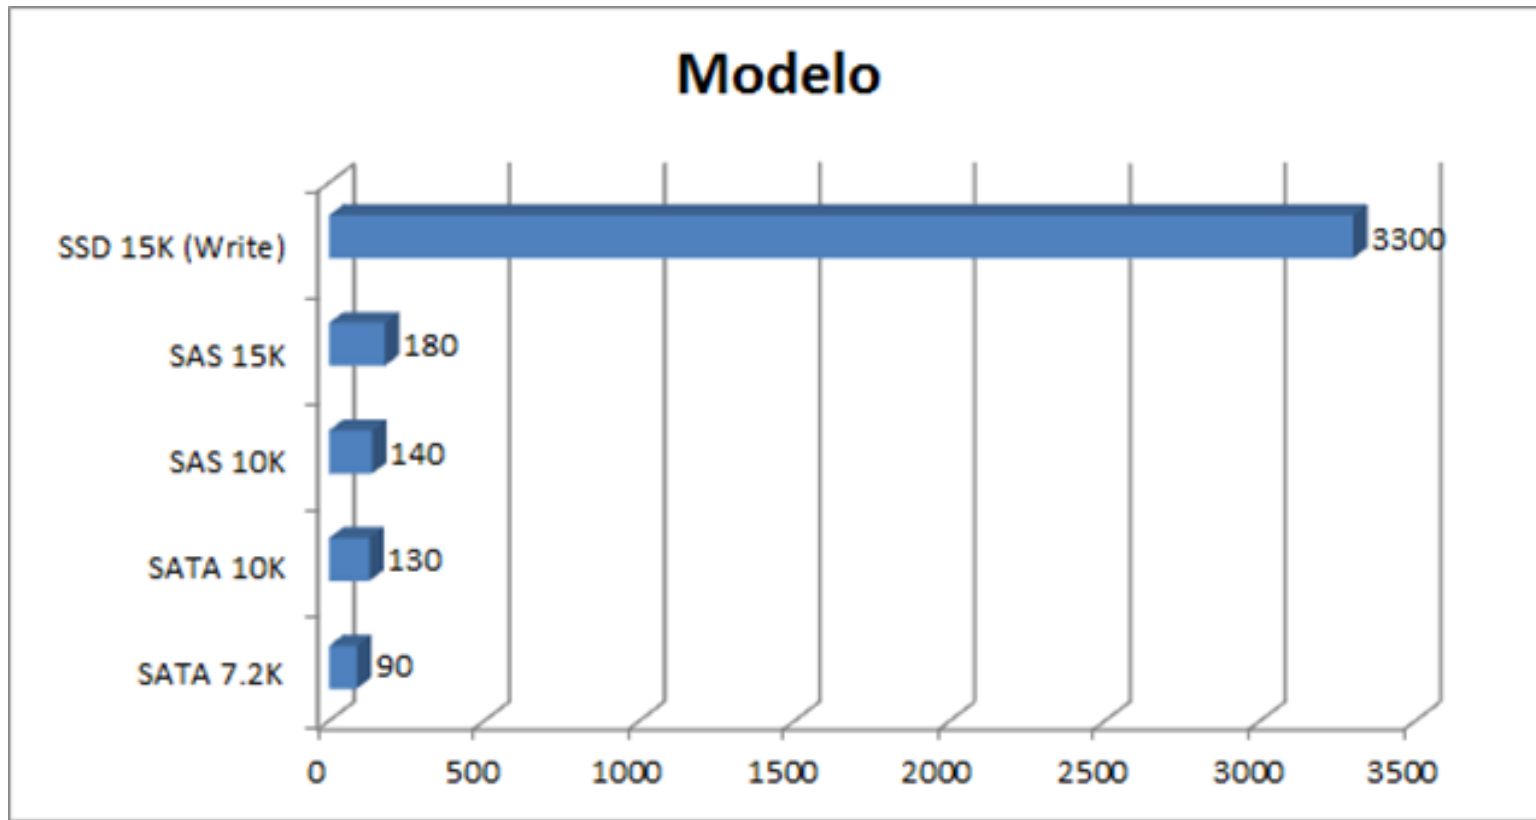
IOPS para SharePoint 2010

Service application	Size estimation recommendation
Search	<p>Search requires three databases. Your environment may include multiple Property and Crawl databases. The Search administration database is typically small: allocate 10 GB.</p> <p>To estimate the required storage for your Property and Crawl databases, use the following multipliers:</p> <ul style="list-style-type: none">• Crawl: $0.046 \times$ (sum of content databases)• Property: $0.015 \times$ (sum of content databases) <p>The IOPS requirements for Search are significant.</p> <ul style="list-style-type: none">• For the Crawl database, search requires from 3,500 to 7,000 IOPS.• For the Property database, search requires 2,000 IOPS.

IOPS para Visio Services


Service Application	Web server CPU	Web server RAM	Application server CPU	Application server RAM	SQL Server CPU	SQL Server IOPS	SQL Server storage
SharePoint Foundation Service	XXX	XXX			XX	XXX	XXX
Central Administration Service			XX	XX	X	X	X
	DELL PE 2950		DELL PE 2950		DELL PE R900		
Role	WFE		Application server		SQL Server-based server		
Processor (CPU)	2pX4 (Xeon L5420@2.5GHz)		2pX4 (Xeon L5420@2.5GHz)		4pX4 (Xeon E7330@2.4GHz)		
RAM in gigabytes (GB)	16		16		32		
Operating system	Windows Server 2008 R2 Enterprise		Windows Server 2008 R2 Enterprise		Windows Server 2008 R2 Datacenter		
Authentication	NTLM		NTLM		NTLM		
Storage: Operating System	4x 146 GB, 10 K RPM, RAID 0		4x 146 GB, 10 K RPM, RAID 0		2x 146 GB, 15 K RPM, RAID 1		
Storage: Backups	--		--		3x 300 GB, 15 K RPM, RAID 5		
Storage: SQL Server data	--		--		9x 300 GB, 15 K RPM, RAID 5		
Storage: SQL Server logs	--		--		6x 300 GB, 15 K RPM, RAID 5		
Number of instances of SQL Server	0		0		1, SQL Server 2008 SP1 CU6		
Business Connection Service *	XX	XX	XXX		XXX		
InfoPath Forms Service	XX	XX	XX		XX	X	X

Tipos de Discos e IOPS



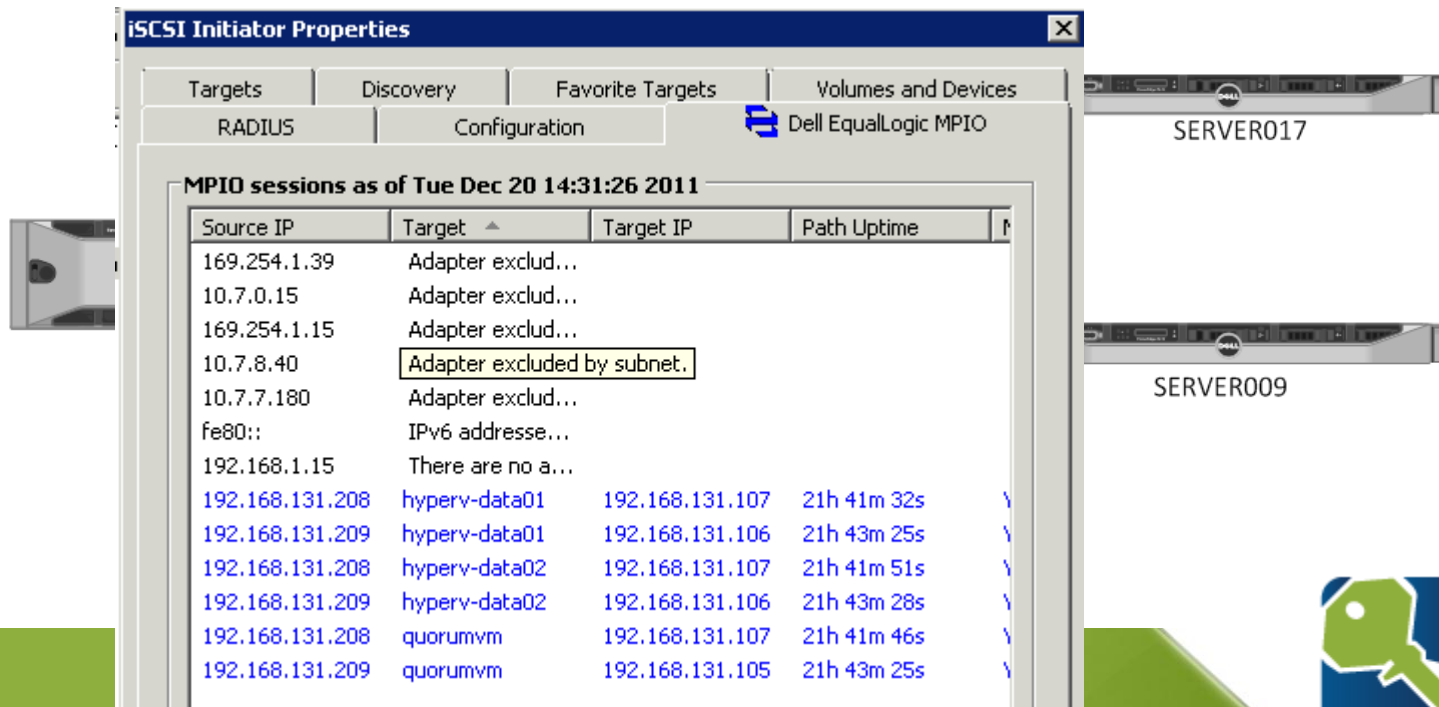
RAID Storage

- RAID Storage é uma CPU?



Multi Path

- Diversas placas acessando o storage, redundância e performance



The screenshot shows the 'iSCSI Initiator Properties' window with the 'Volumes and Devices' tab selected. The 'Dell EqualLogic MPIO' configuration is active. The 'MPIO sessions as of Tue Dec 20 14:31:26 2011' table lists various source IP addresses and their corresponding target information. The 'Adapter excluded by subnet' entry is highlighted.

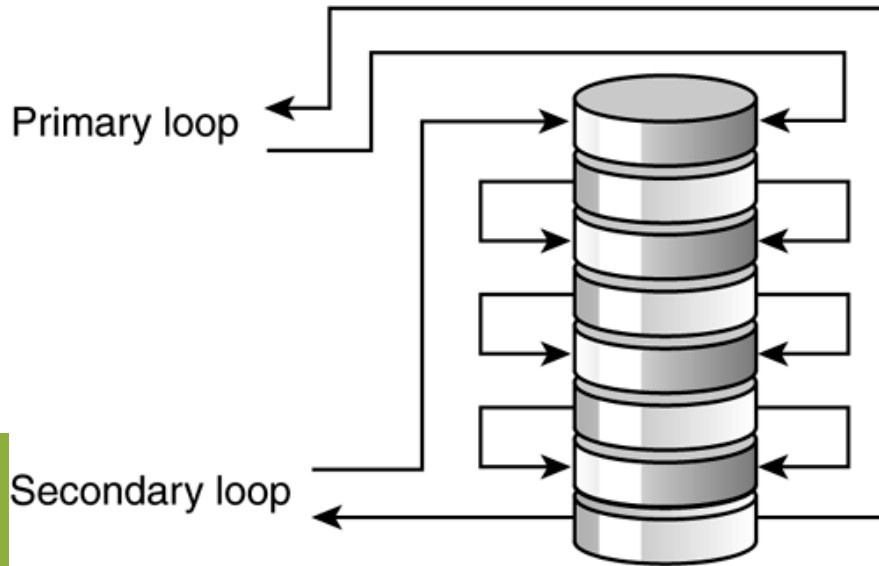
Source IP	Target	Target IP	Path Uptime
169.254.1.39	Adapter exclud...		
10.7.0.15	Adapter exclud...		
169.254.1.15	Adapter exclud...		
10.7.8.40	Adapter excluded by subnet.		
10.7.7.180	Adapter exclud...		
fe80::	IPv6 address...		
192.168.1.15	There are no a...		
192.168.131.208	hyperv-data01	192.168.131.107	21h 41m 32s
192.168.131.209	hyperv-data01	192.168.131.106	21h 43m 25s
192.168.131.208	hyperv-data02	192.168.131.107	21h 41m 51s
192.168.131.209	hyperv-data02	192.168.131.106	21h 43m 28s
192.168.131.208	quorumvm	192.168.131.107	21h 41m 46s
192.168.131.209	quorumvm	192.168.131.105	21h 43m 25s

Configurações de Storage

- JBOD
- RAID Storage
- RAID 0
- RAID 1
- RAID 5 e 50
- RAID 6
- RAID 10

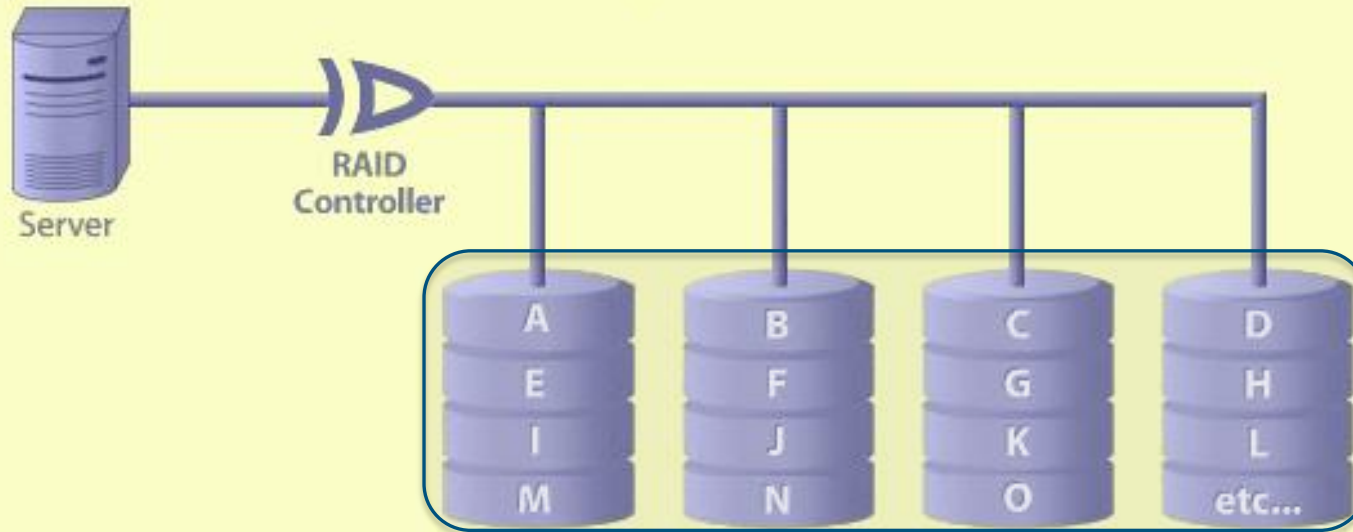
JBOD

“Common type of disk storage where multiple disks are attached to a common motherboard (backplane) and appear as individual attached devices to the network (SAN)”.



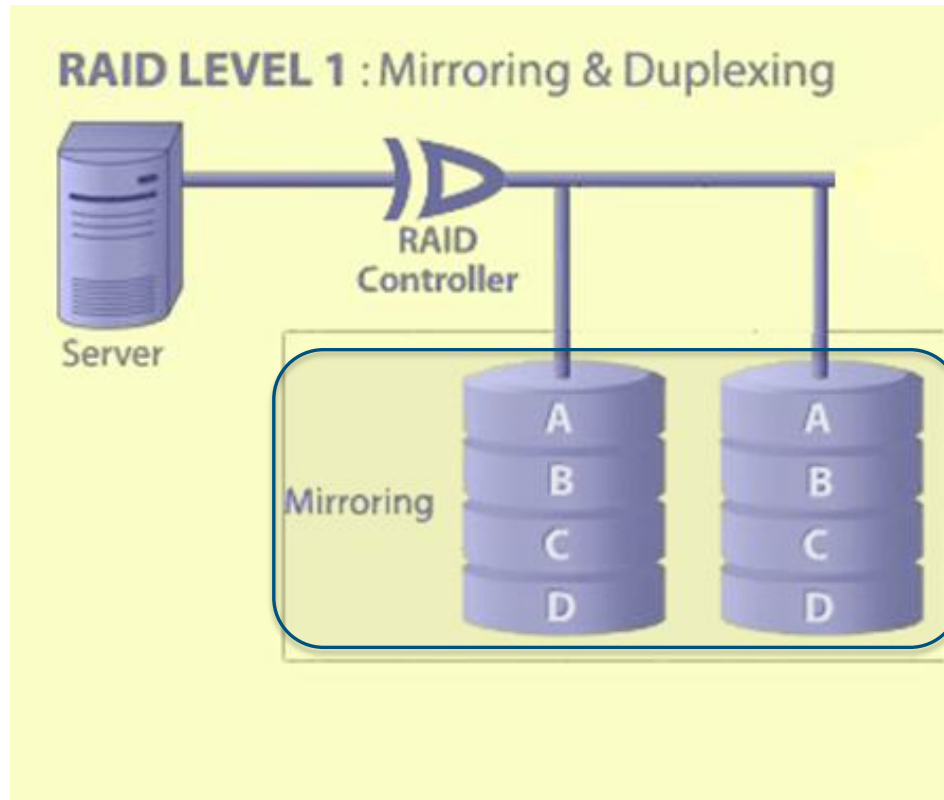
RAID 0

RAID LEVEL 0 : Striped Disk Array without Fault Tolerance



- RAID 0 requer mínimo de 2 discos
- Não tem redundância, mas boa performance e utilização de espaço


RAID 1



- RAID 1 requer mínimo de 2 discos e trabalha em pares
- Ótima redundância e performance mas péssima utilização de espaço
- Recomenda-se utilizar controladoras separadas quando sem enclosures

RAID 5

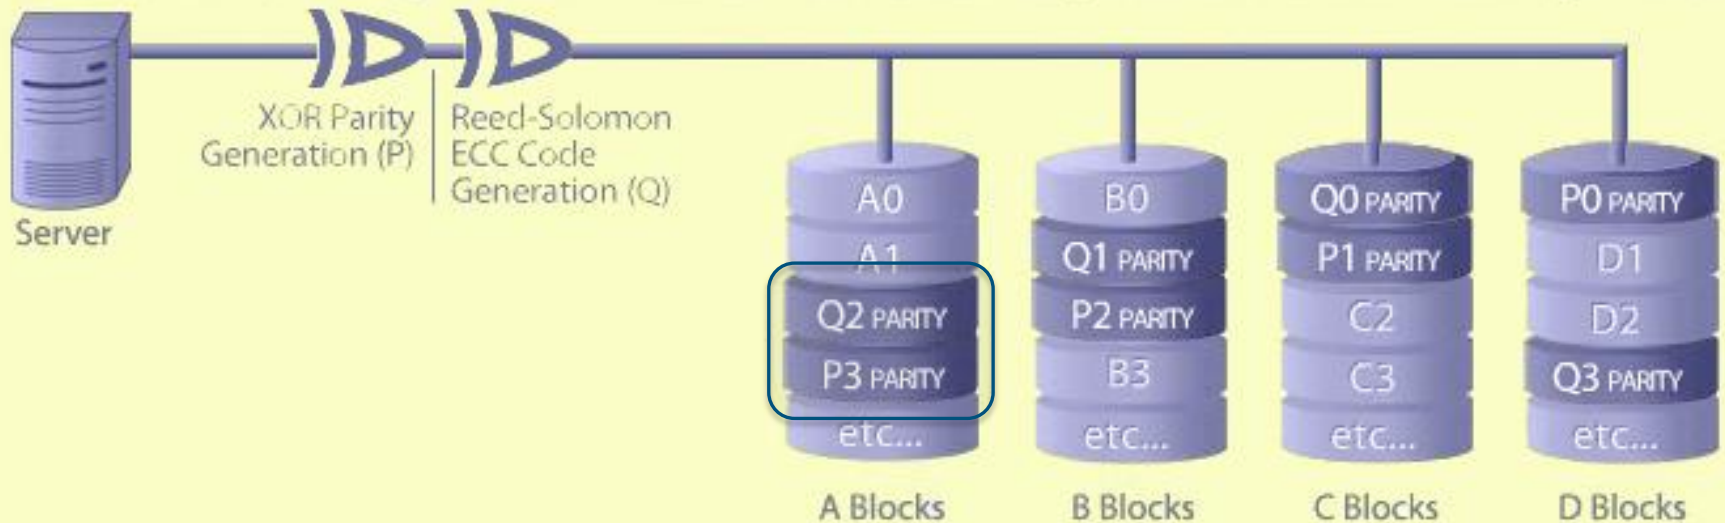
RAID LEVEL 5 : Independent Data Disks with Distributed Parity Blocks



- RAID 5 requer mínimo de 3 discos, um para paridade
- Boa redundância, performance e utilização do espaço
- RAID 50 duplica a paridade em um disco (spare), exigindo 4 discos mínimo, aumentando a redundância

RAID 6

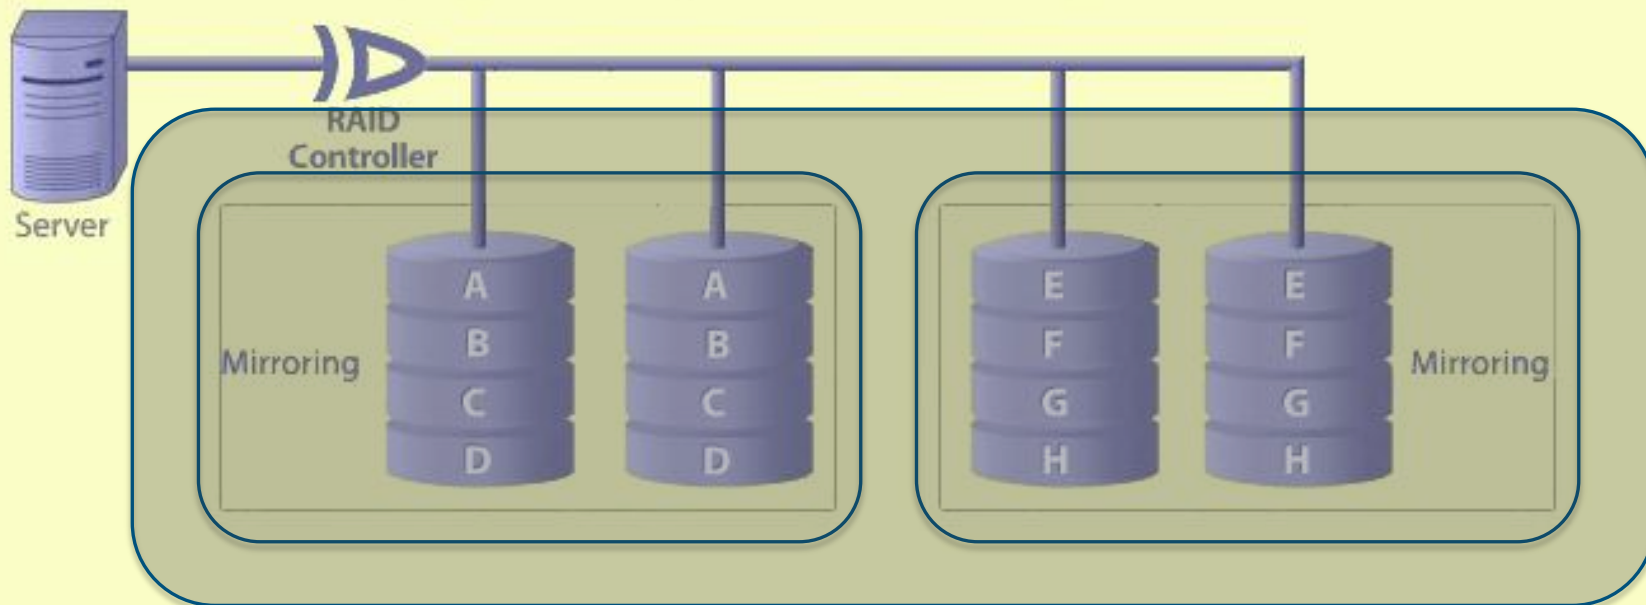
RAID LEVEL 6 : Independent Data Disks with Two Independent Distributed Parity Schemes



- RAID 6 requer mínimo de 4 discos, dois para paridade
- Ótima redundância, boa performance e utilização do espaço ideal para muitas cabeças

RAID 10

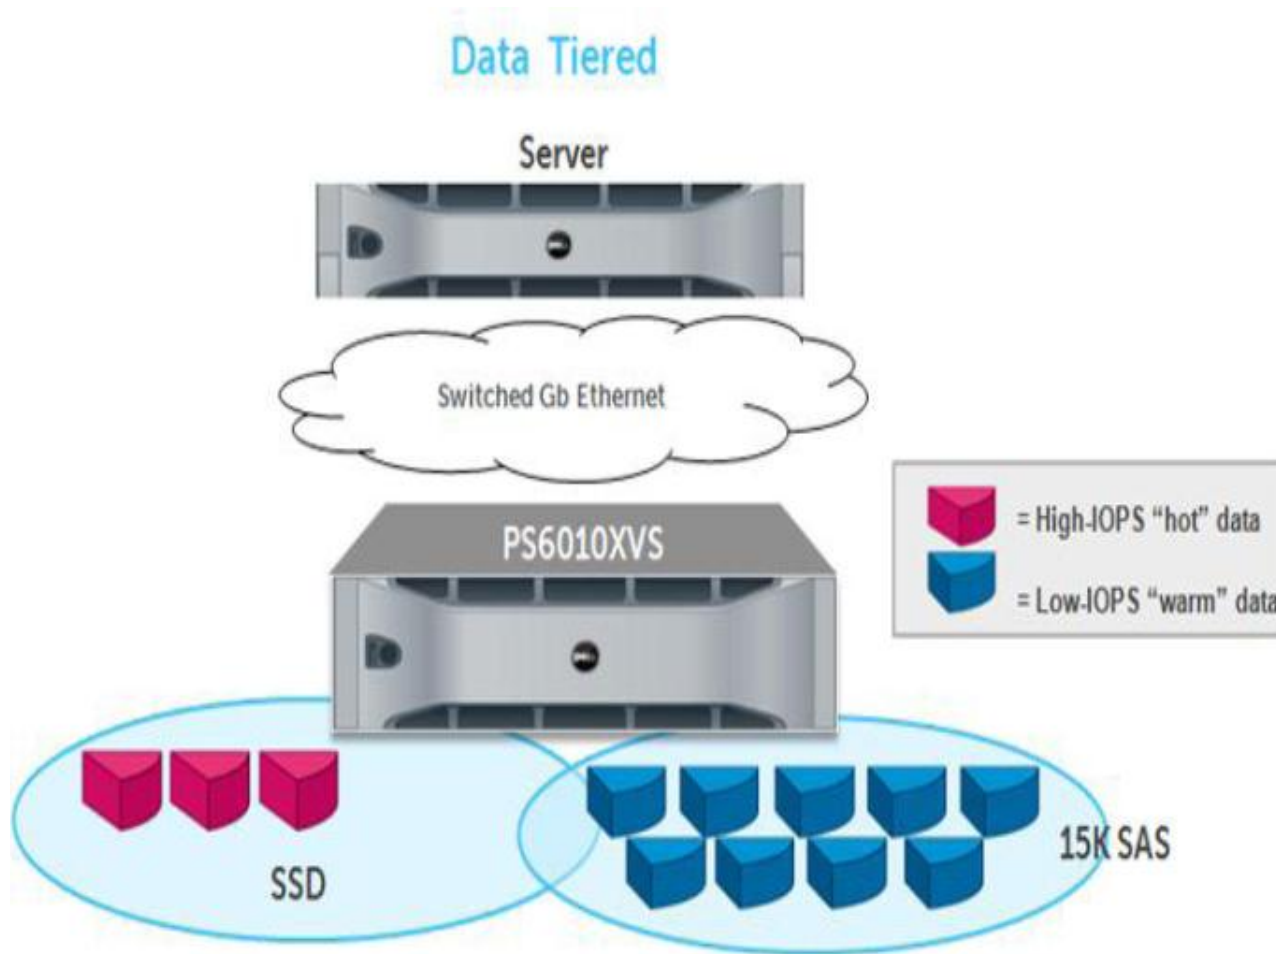
RAID LEVEL 10 : Very High Reliability Combined with High Performance



- RAID 10 requer mínimo de 4 discos
- Ótima redundância e performance, péssima utilização do espaço
- É utilizado com controladoras separadas
- Melhor que RAID 0+1

Hybrid Arrays (or Enclosures)

- AUTOMATED WORKLOAD TIERING**



IOPS do SQL Server com SAP

I/O type	Average IOPs	Sustained maximum IOPs	Peak IOPs	Sustained maximum MB/s	Peak MB/s	Average I/O size
Data file reads	2,500 (over a 24-hour period) 4,500 (during busy times)	10,000	15,000	400 MB/s	440 MB/s	20K
Data file writes	Less than 200	1,000	1,900	70 MB/s	85 MB/s	25K
Log file writes	Less than 200	200	200	40 MB/s	16 MB/s	16K (60K maximum)

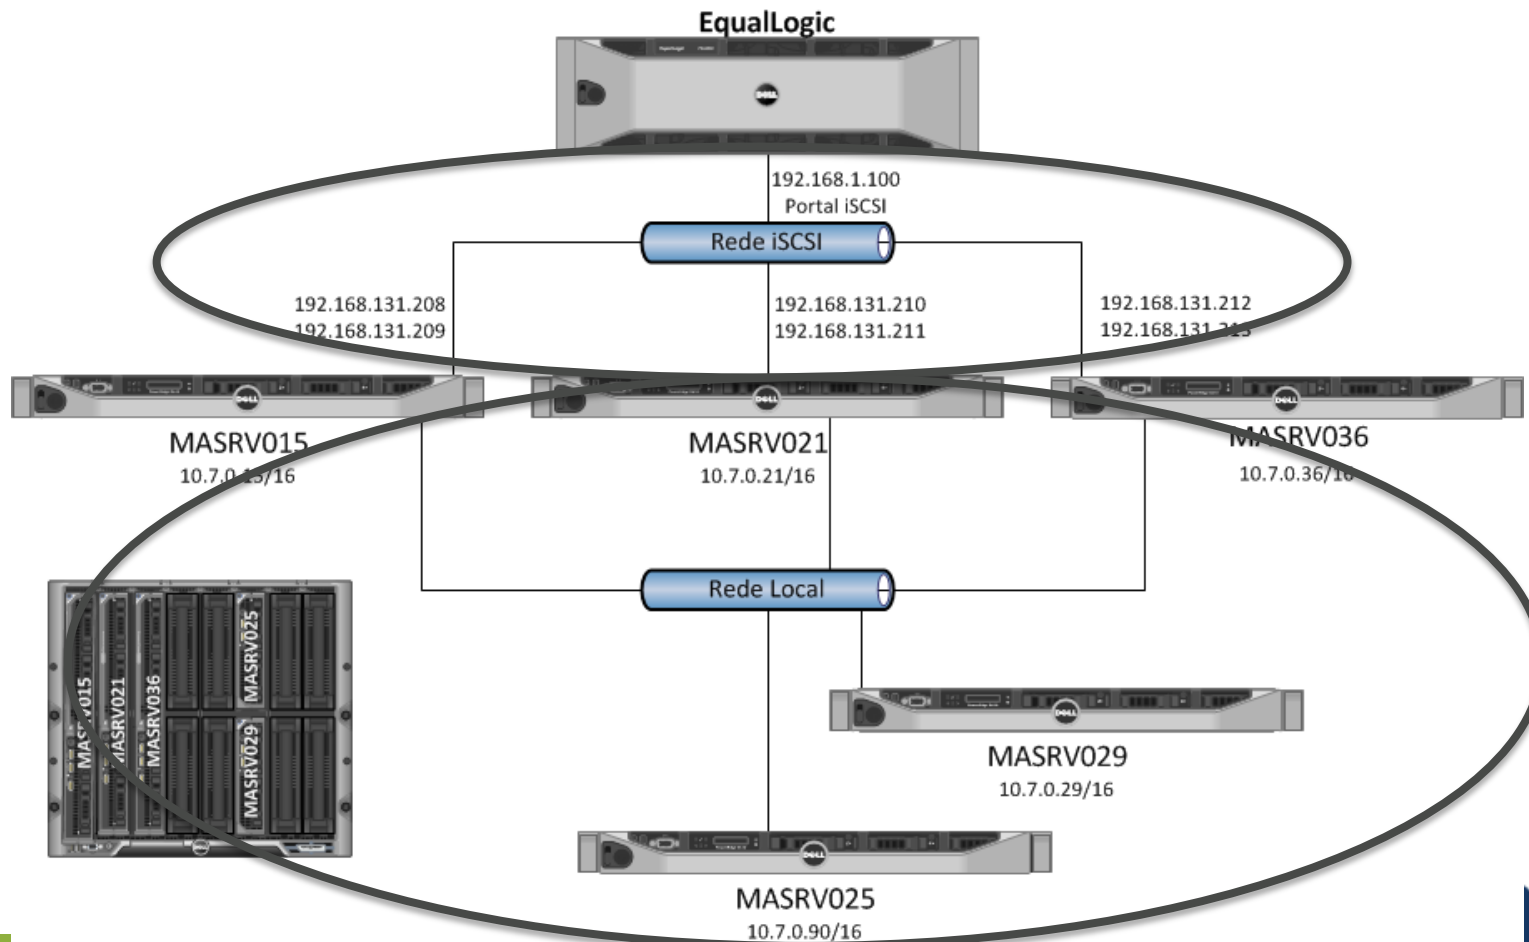
Disks	Number of disks	RAID level	Total logical I/Os (as measured via <u>perfmon</u>)	Total physical I/O (RAID adjusted)	I/Os per device
SSD	4	5	12,044	17,876	4,469
Traditional 15K <u>Fibre Channel</u> drives	34	1+0	12,362	14,465	425

Network

- Como utilizar os recursos do hardware de forma a ter a melhor performance da rede do cliente

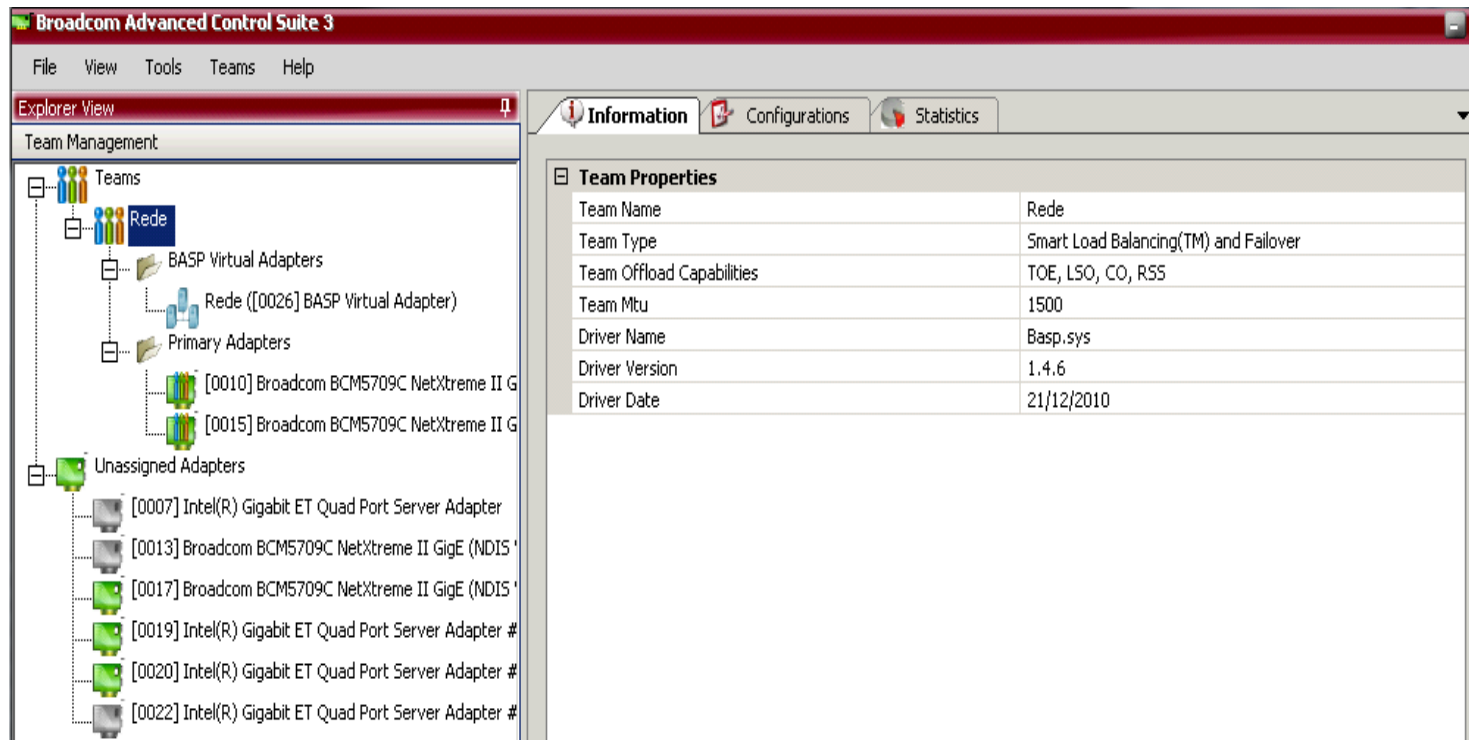
Redes Físicas Segregadas

Diagrama de Rede dos Servidores



NIC Team

- Multiplos adaptadores acessando a mesma rede



The screenshot displays the Broadcom Advanced Control Suite 3 interface. The left pane shows a tree view under 'Team Management' with a team named 'Rede'. This team is composed of 'BASP Virtual Adapters' (one 'Rede ([0026] BASP Virtual Adapter)') and 'Primary Adapters' (two Broadcom BCM5709C NetXtreme II Gigabit adapters with IDs [0010] and [0015]). Below this, a list of 'Unassigned Adapters' includes several Intel and Broadcom network cards.

The right pane shows the 'Team Properties' for the 'Rede' team:

Team Properties	
Team Name	Rede
Team Type	Smart Load Balancing(TM) and Failover
Team Offload Capabilities	TOE, LSO, CO, RSS
Team Mtu	1500
Driver Name	Basp.sys
Driver Version	1.4.6
Driver Date	21/12/2010

Processador

- Qual o melhor tipo e configuração de processamento para ambientes de banco de dados

Processador por Versão

SQL Server Edition	Maximum Compute Capacity Used by a Single Instance (SQL Server Database Engine)	Maximum Compute Capacity Used by a Single Instance (AS, RS)
Enterprise Edition: Core-based Licensing ¹	Operating system maximum	Operating system maximum
Developer	Operating system maximum	Operating system maximum
Evaluation	Operating system maximum	Operating system maximum
Business Intelligence	Limited to lesser of 4 Sockets or 16 cores	Operating system maximum
Standard	Limited to lesser of 4 Sockets or 16 cores	Limited to lesser of 4 Sockets or 16 cores
Web	Limited to lesser of 4 Sockets or 16 cores	Limited to lesser of 4 Sockets or 16 cores
Express	Limited to lesser of 1 Socket or 4 cores	Limited to lesser of 1 Socket or 4 cores
Express with Tools	Limited to lesser of 1 Socket or 4 cores	Limited to lesser of 1 Socket or 4 cores
Express with Advanced Services	Limited to lesser of 1 Socket or 4 cores	Limited to lesser of 1 Socket or 4 cores

Configuração de Cores/Sockets

- Servidores Fisicos
 - Utilizar HyperThreading e respeitar os limites da versão do SQL Server
- Servidores Virtuais
 - Não utilizar HyperTheading, pois o Hypervisor apresenta VPs
- Configure corretamente o “Max Degree Paralellism”

Patrocinadores



O'REILLY®

