



Desempenho dos nós de processamento do cluster SEARCH

Vítor Oliveira

Universidade do Minho

Nov-05



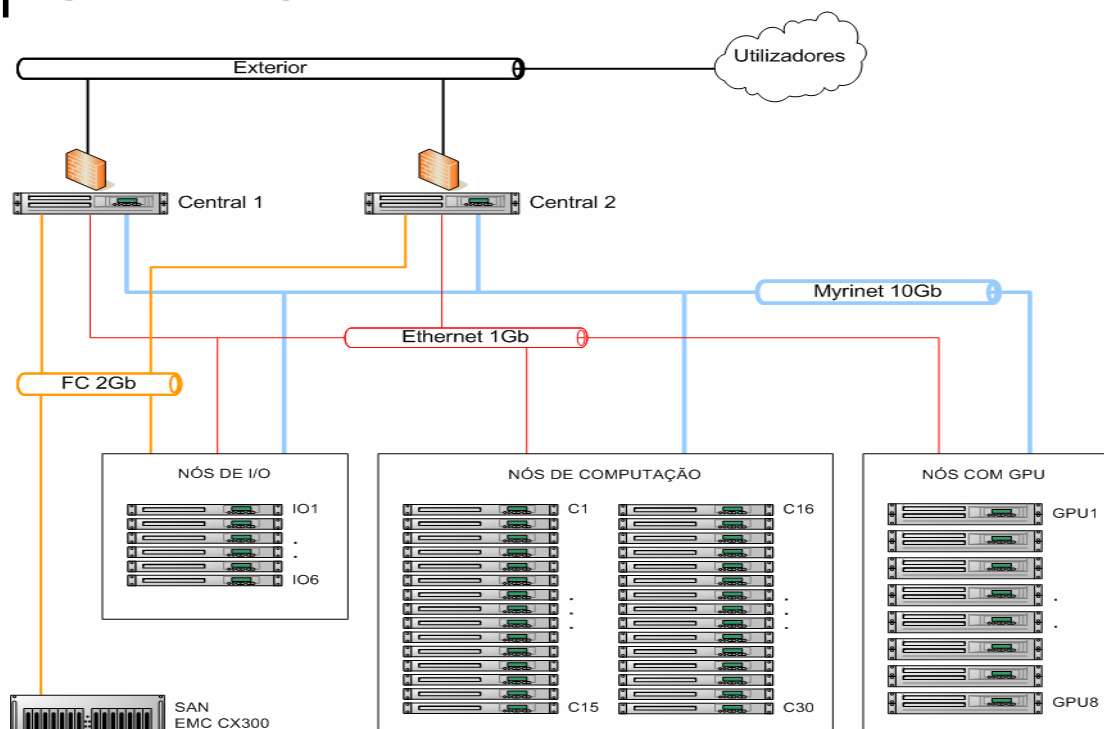
Parte 1: Apresentação



Alguns números do cluster SEARCH

- 46 nós de computação: 36 de 1U de altura e 8 de 2U de altura
- 92 processadores Xeon a 3.2 GHz (588,8 GFlops max)
- 92 GB de RAM
- 8 processadores NVidia 7800GTX (300M transístores)
- 920 Gbps Myrinet-10G
- 92 Gbps Gb Ethernet
- 3,0 TB de armazenamento em SAN
- 3,7 TB de armazenamento em disco local
- 17,5 kW de consumo de energia eléctrica
- 59,5 kBTU/h de calor gerado
- 1,5 toneladas

Estrutura do cluster SEARCH



Nós de computação SR1435VP2



Vista aérea da motherboard SR7320VP2

Ref #	Description	Ref #	Description
1	(J1A1) 2-pin Chassis Intrusion Header (J1A2) 2-pin Hard Drive Act LED Header (J1A4) Rolling BIOS Jumper	I	CPU #1 Fan Header
2	10-pin DH10 Serial A Header	J	5-pin Power Sense Header
3	USB Port 2	K	CPU #2 Socket
4	USB Port 1	L	CPU #1 Socket
5	Video Connector	M	6300ESB ICH – Chipset Component
6	NIC #2	N	SATA Ports
7	NIC #1	O	(J1H2) Password Clear Jumper (J1H3) Recovery Boot Jumper (J1H5) CMOS Clear Jumper
8	RJ45 Serial B Port	P	Legacy ATA-100 connector
9	Stacked PS/2 Keyboard and Mouse Ports	Q	50-pin Control Panel Header
A	CMOS Battery	R	100-pin Control Panel, Floppy, IDE Connector
B	Full-height Riser Card Slot	S	Legacy Floppy Connector
C	Low-profile Riser Card Slot	T	SSI 34-pin Control Panel Header
D	DIMM Slots	U	8-pin AUX Power Connector
E	MCH – Chipset Component	V	24-pin Main Power Connector
F	1x10 USB Header	W	SSI System Fan Header
G	ATI RageXL Video Controller	X	Server Chassis SR1400LC / SR2400 System Fan Header
H	CPU #2 Fan Header	Y	Processor Voltage Regulator Circuitry




Diagrama funcional dos nós de computação do SEARCH

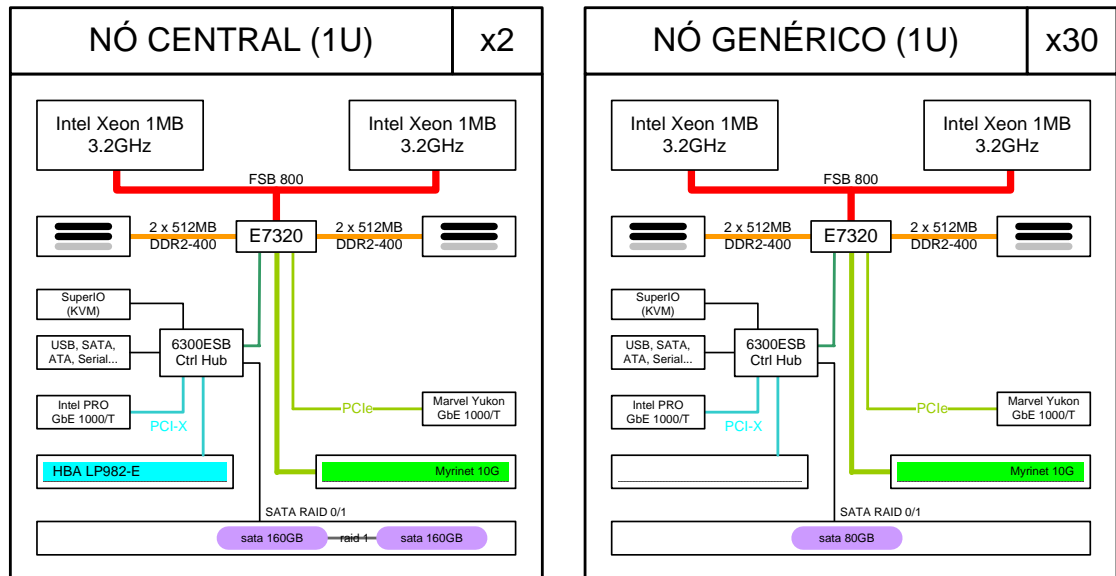


Diagrama funcional dos nós de computação do SEARCH

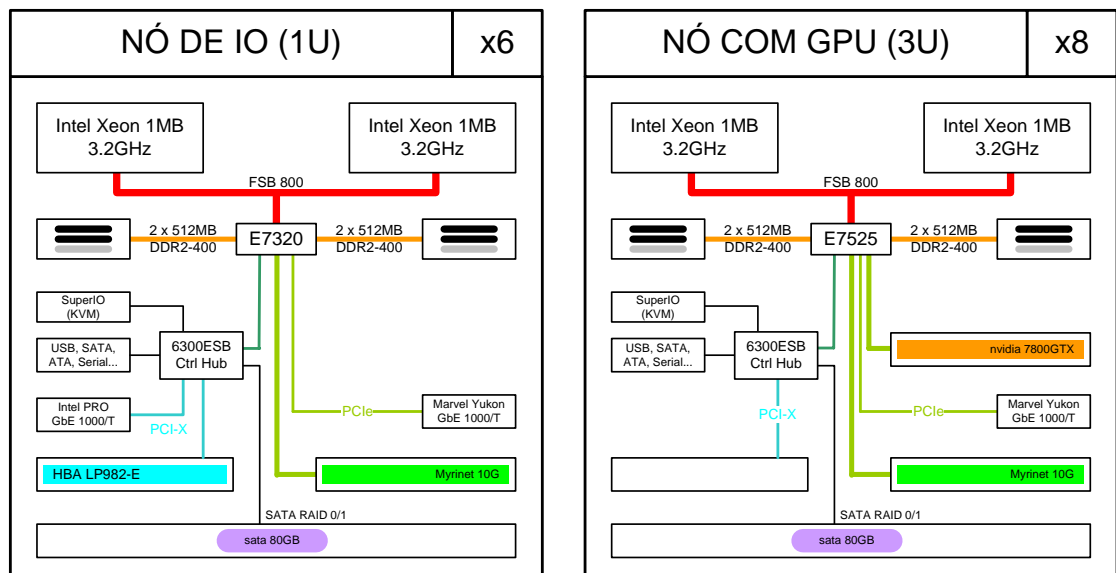


Diagrama funcional da placa mãe SE7320VP2

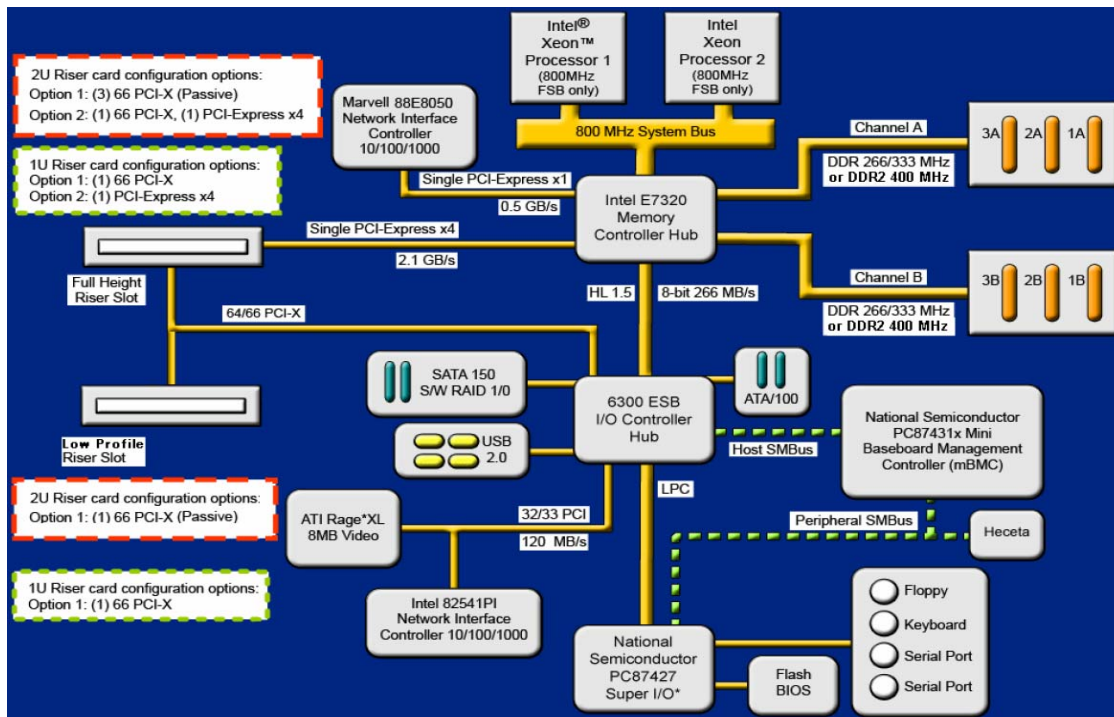
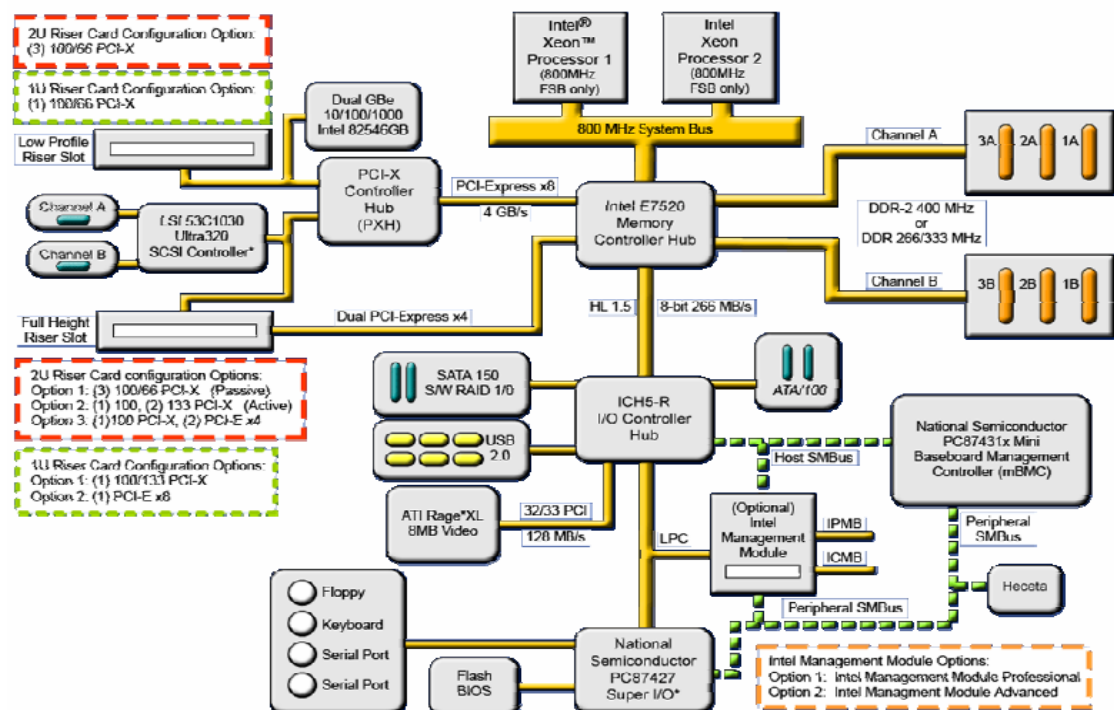


Diagrama funcional da placa mãe SE7520JR2 (alternativa)





Parte 2: O desempenho

Largura de banda

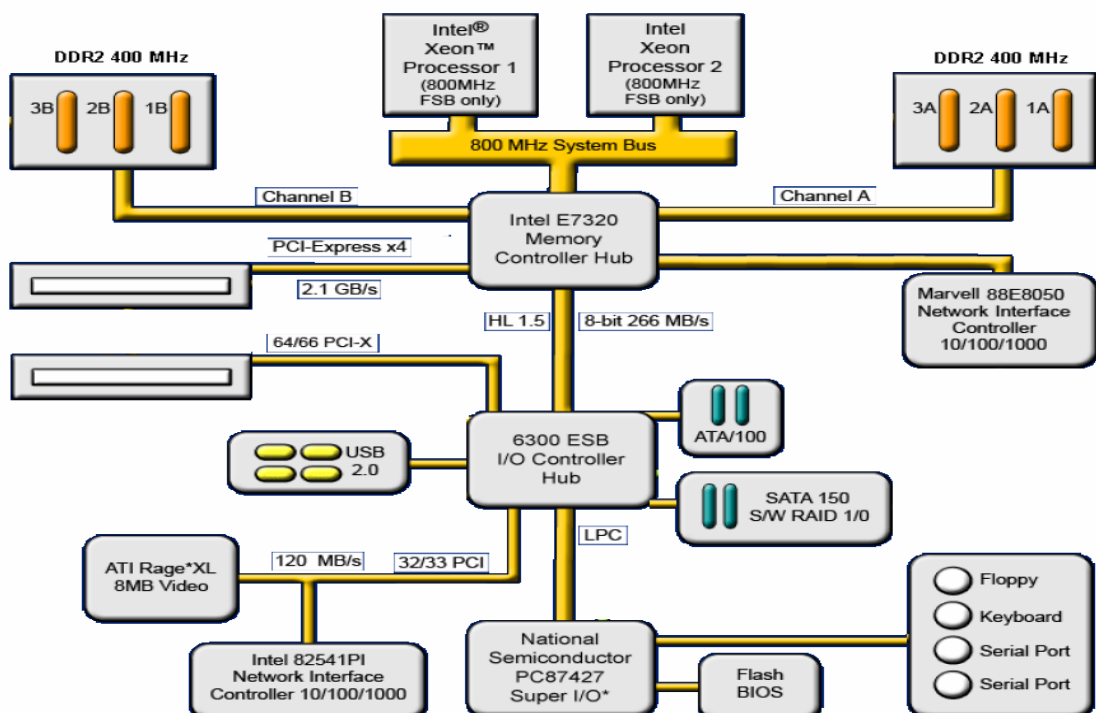
O débito de dados suportado pelos barramentos de comunicações.

Latência

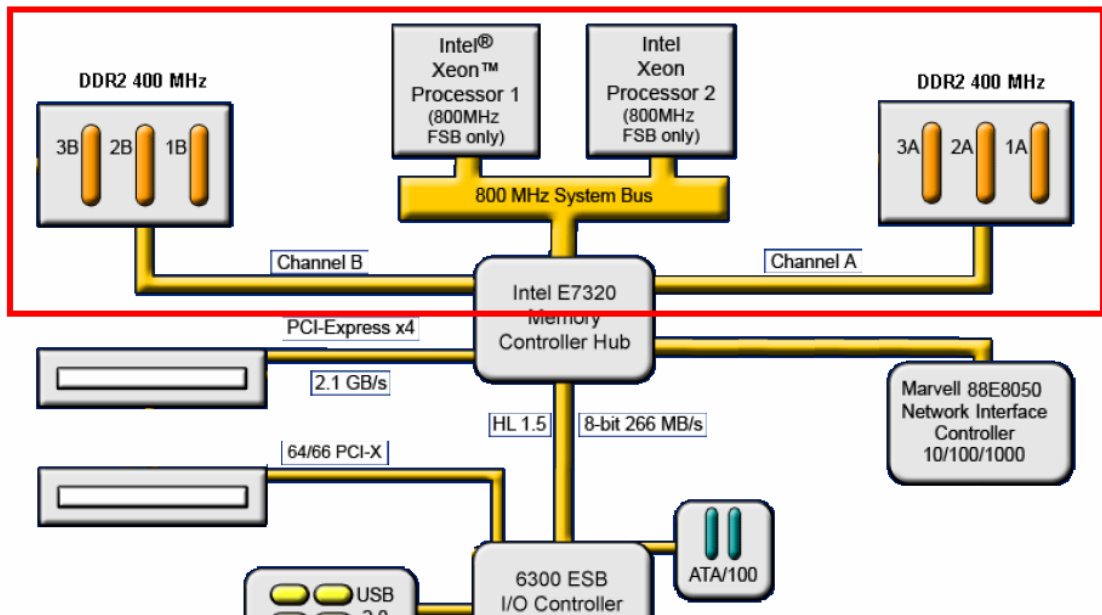
O tempo mínimo necessário para um sinal atravessar os barramentos de comunicações.




Diagrama funcional simplificado



FSB e memória



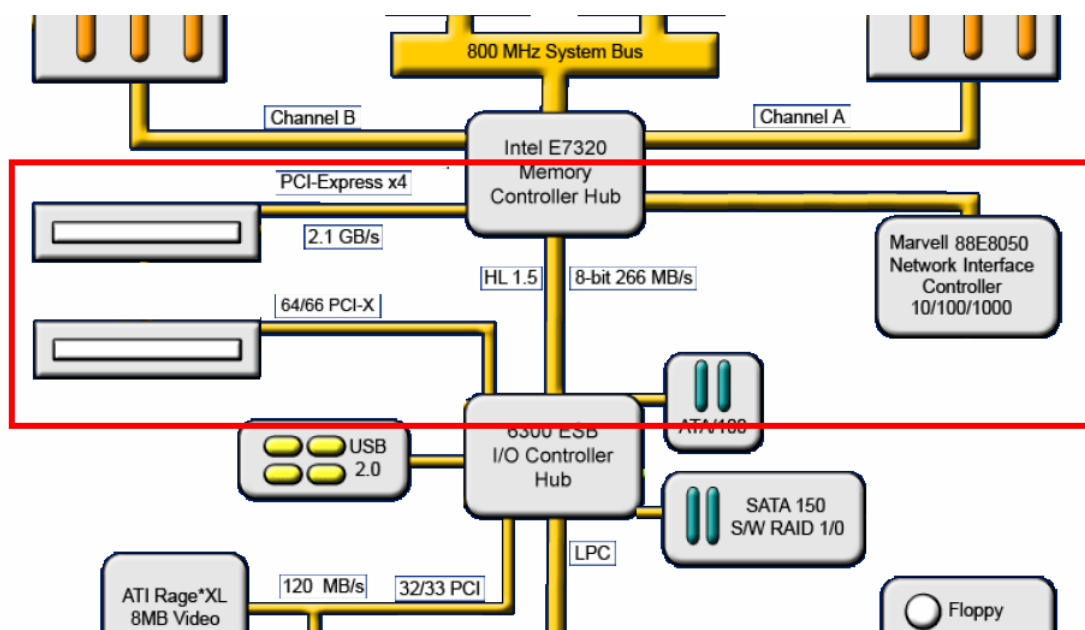
Largura de banda dos barramentos FSB e memória



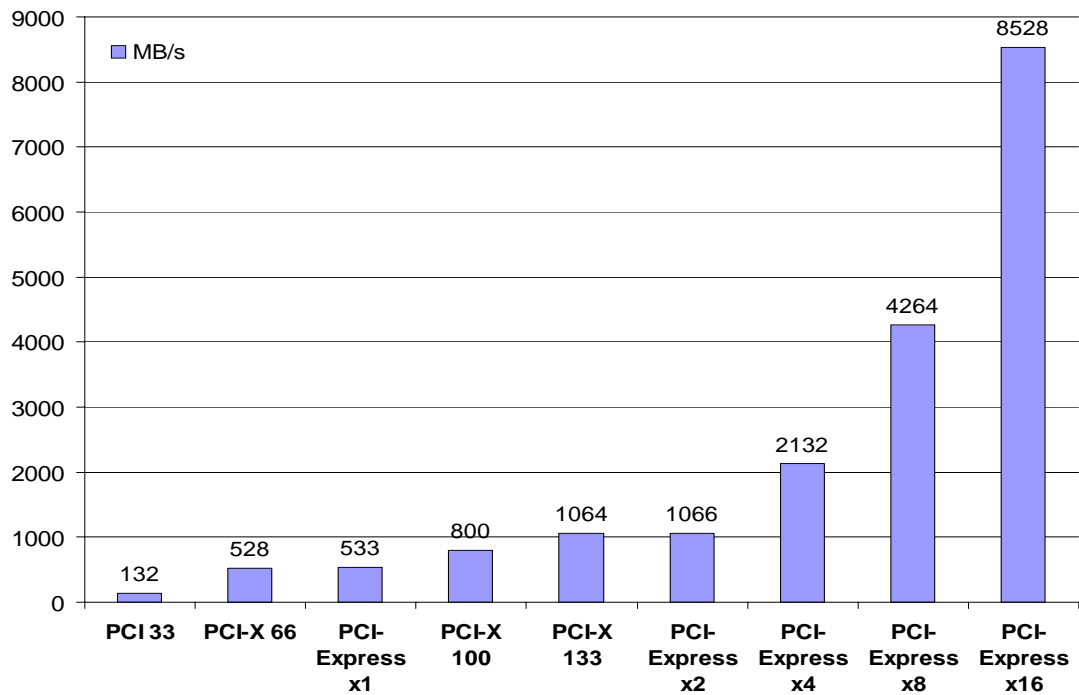
A latência e a velocidade da luz

- $C=300\ 000\ 000\ \text{m} / \text{s}$
- Propagação da electricidade no cobre é 80% da da velocidade da luz
- Considere um CPU com relógio a 3.2GHz ($f=3\ 200\ 000\ 000\text{Hz}$)
- Em um ciclo do relógio o sinal eléctrico percorre $C/f\ (\text{m}) * 80\%$
- => **7,5 cm/ciclo** do relógio do CPU

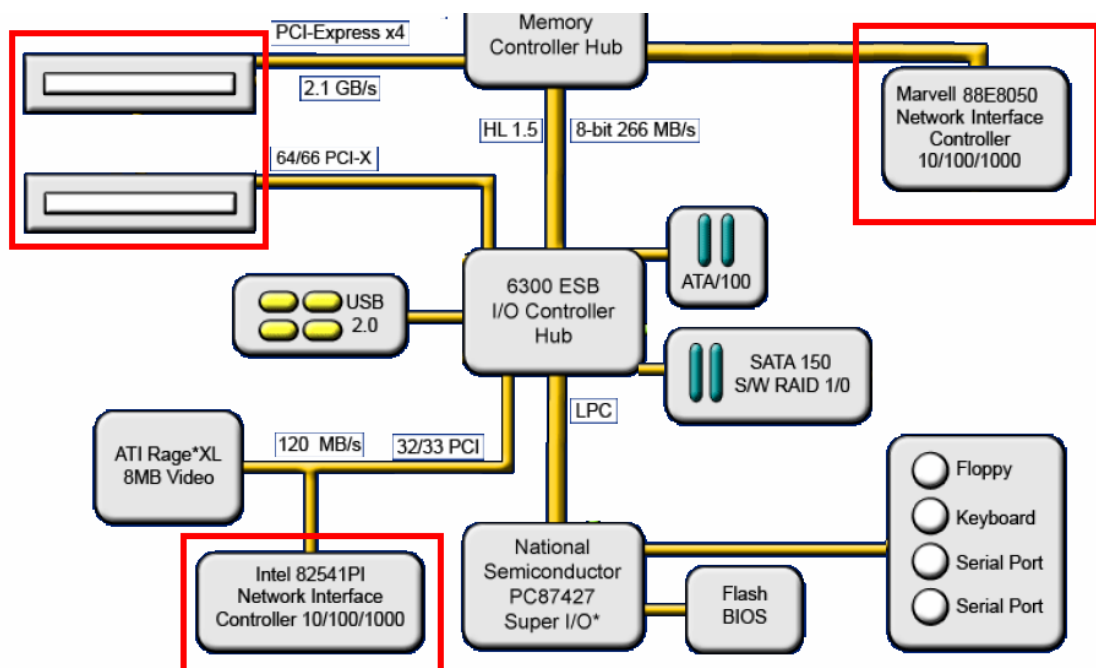
Barramentos PCI



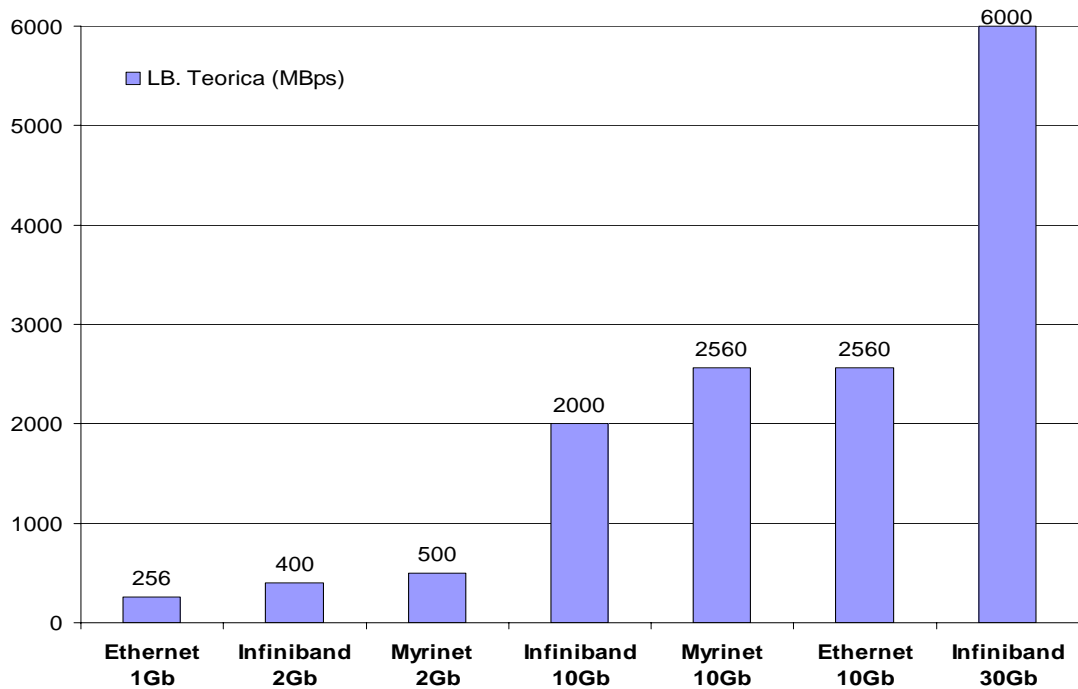
Largura de banda dos barramentos PCI




Comunicações



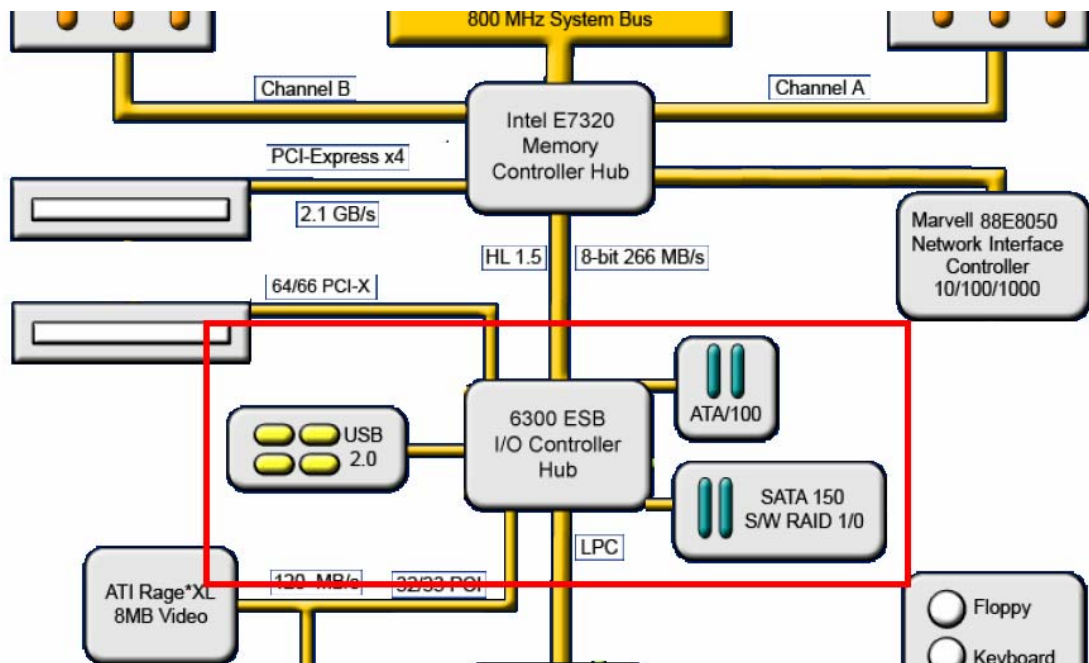
Largura de banda das redes de comunicações



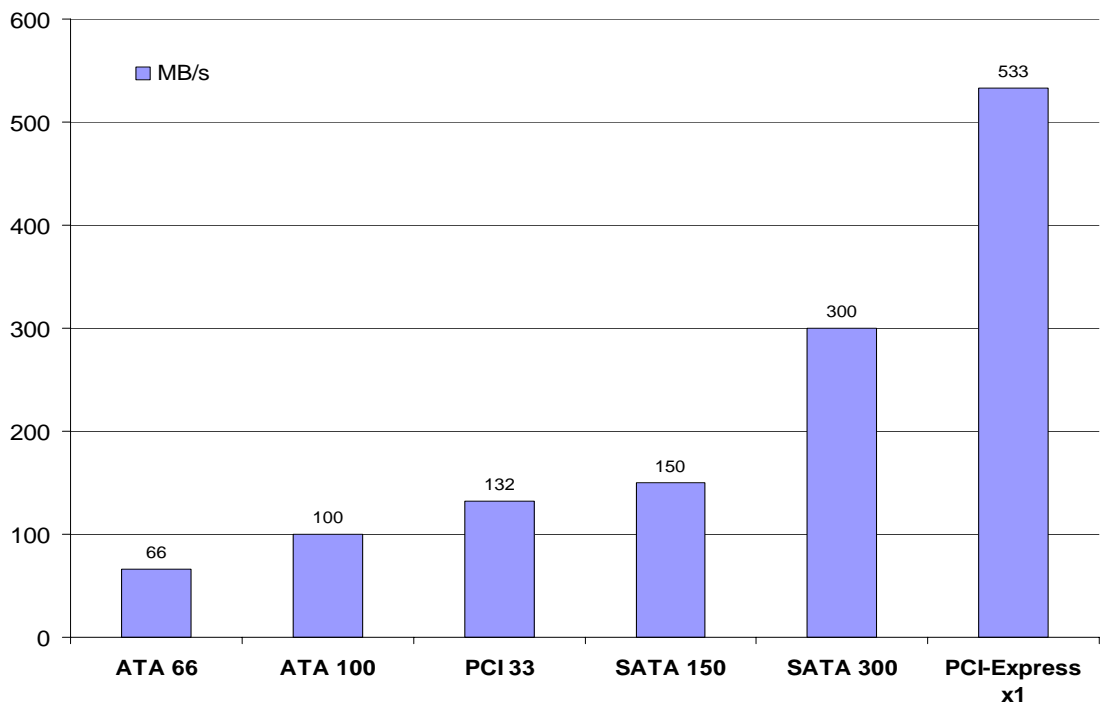
Comparativo da largura de banda dos barramentos



Entrada e saída de dados para disco (I/O)

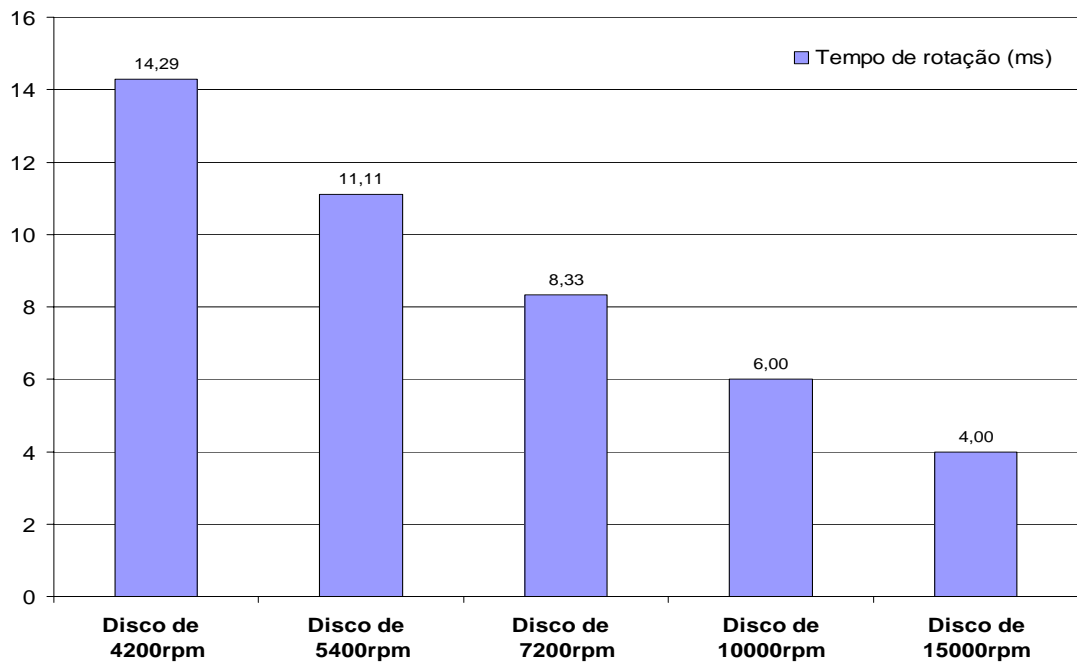


Largura de banda de I/O






Latência de acesso aos discos

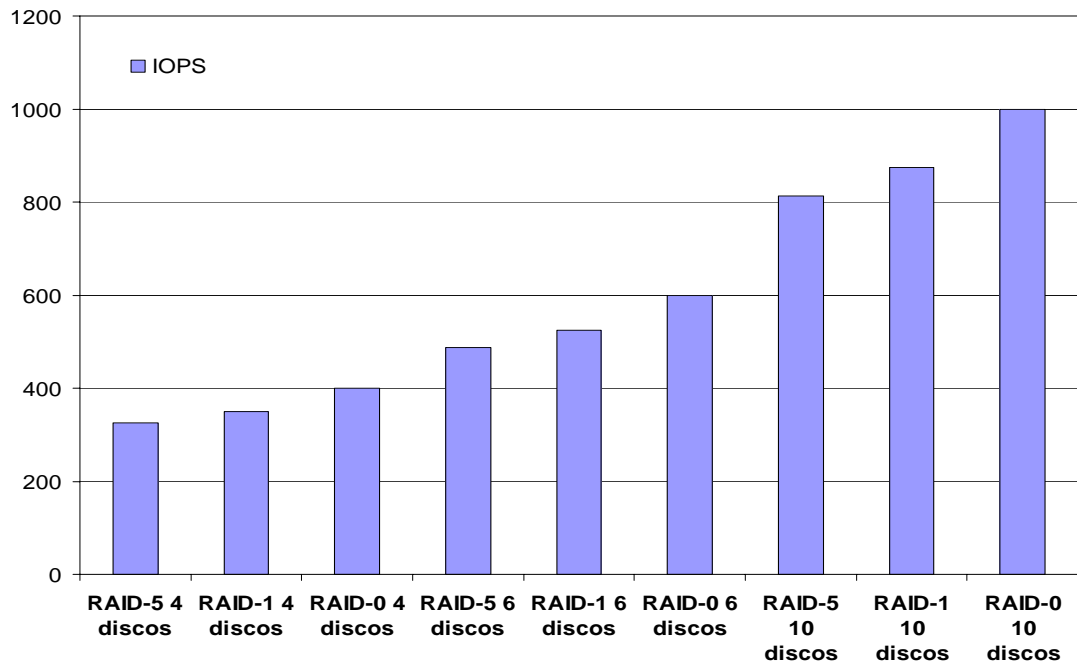


Número de operações de I/O por segundo (IOPS)





Número de operações de I/O por segundo com RAID



Fim

Visita ao cluster SEARCH