Cluster SeARCH

Arquitectura física

Apresentação

Alguns dados do cluster SEARCH (2006)

- 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 trans.)
- 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

Nós de computação

Diagramas funcionais dos nós do cluster do SEARCH

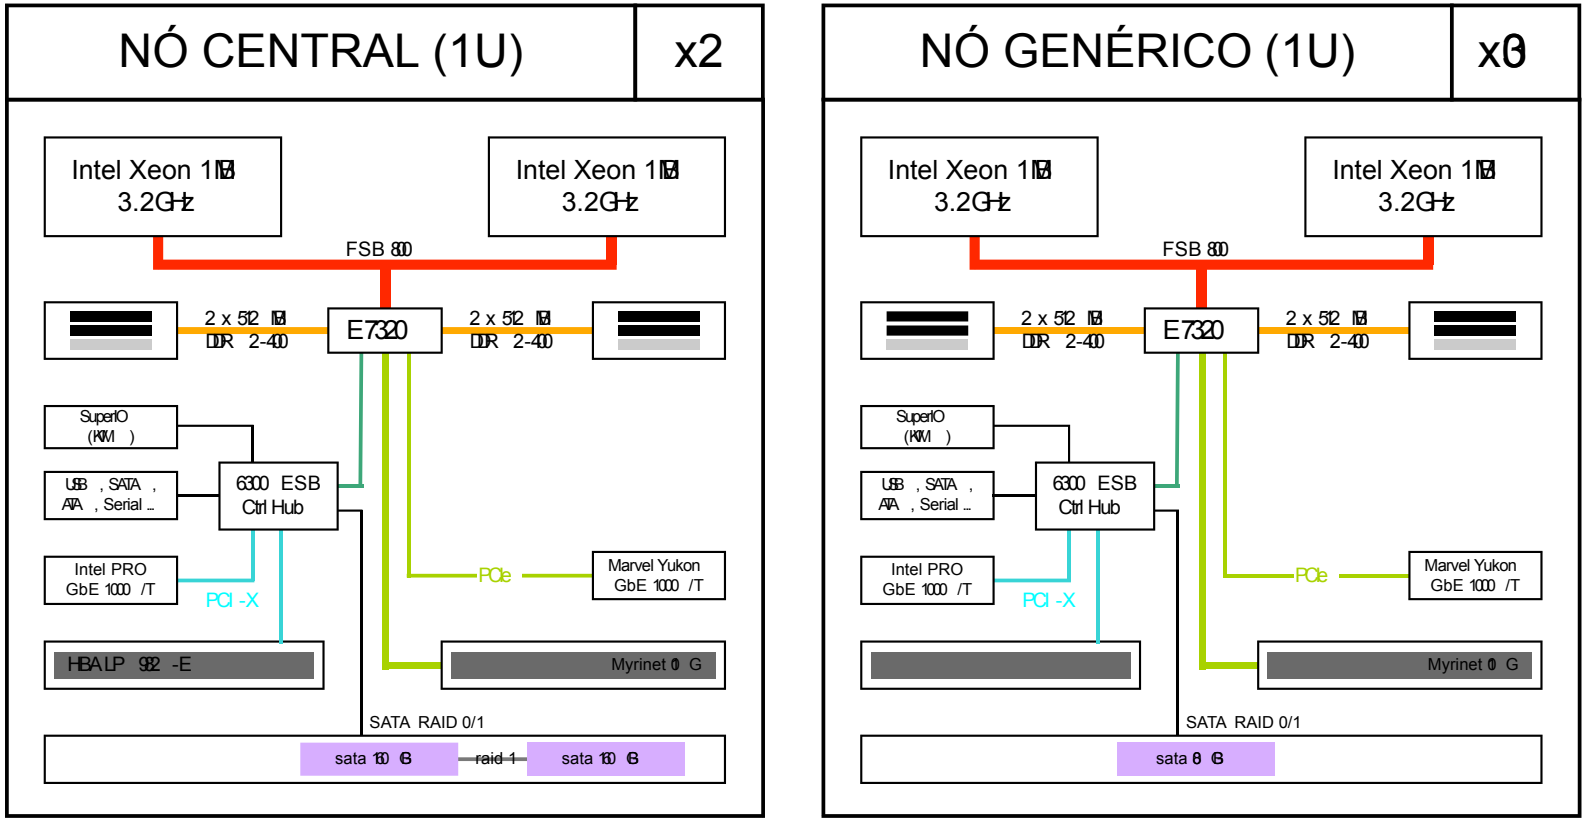

Diagramas funcionais dos nós do cluster do SEARCH

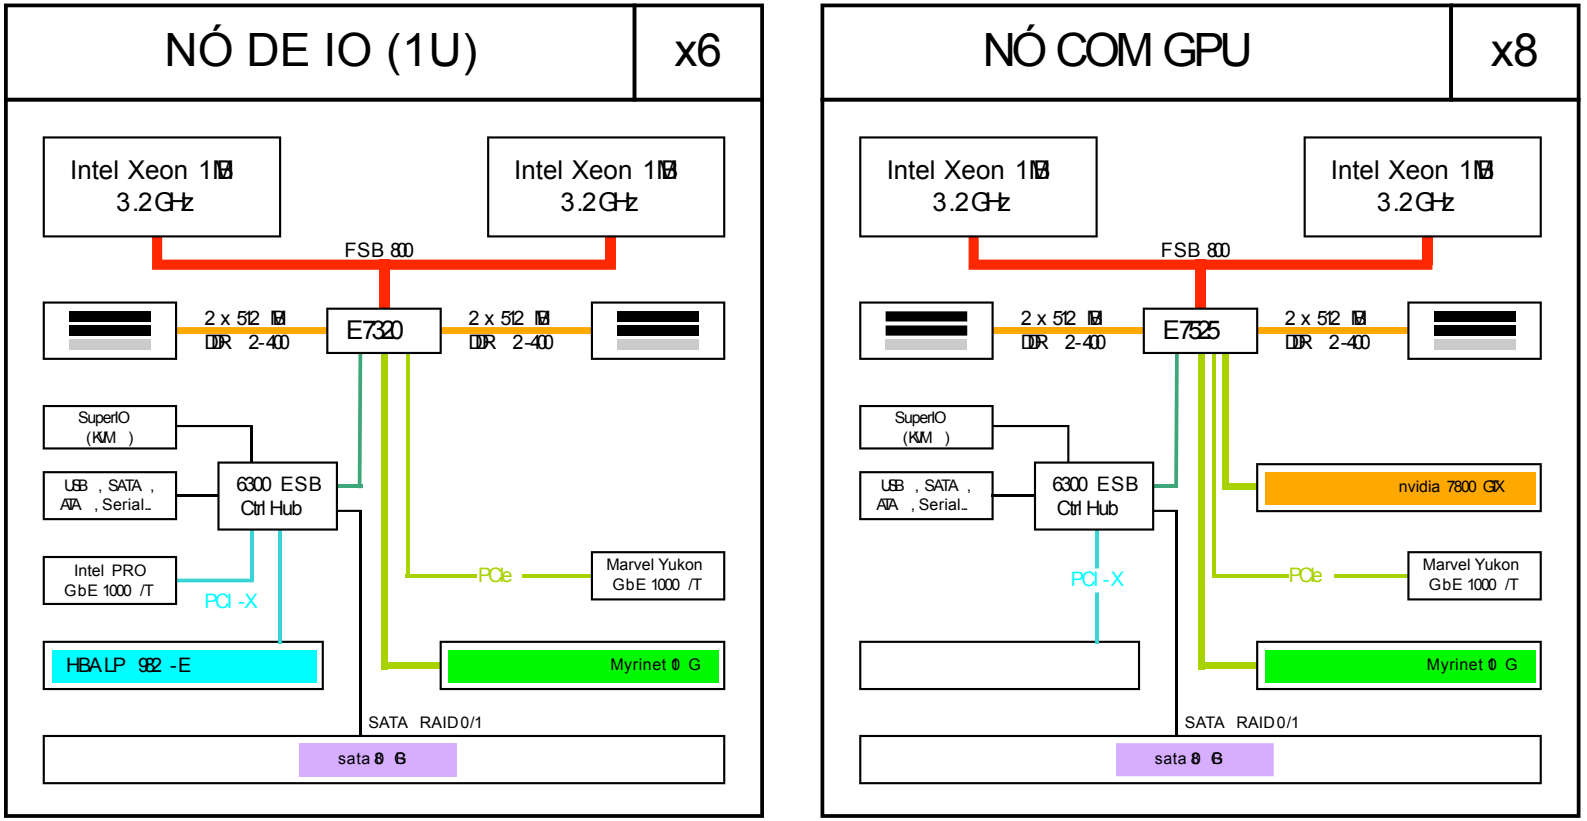

Vista aérea da motherboard SR7320VP2

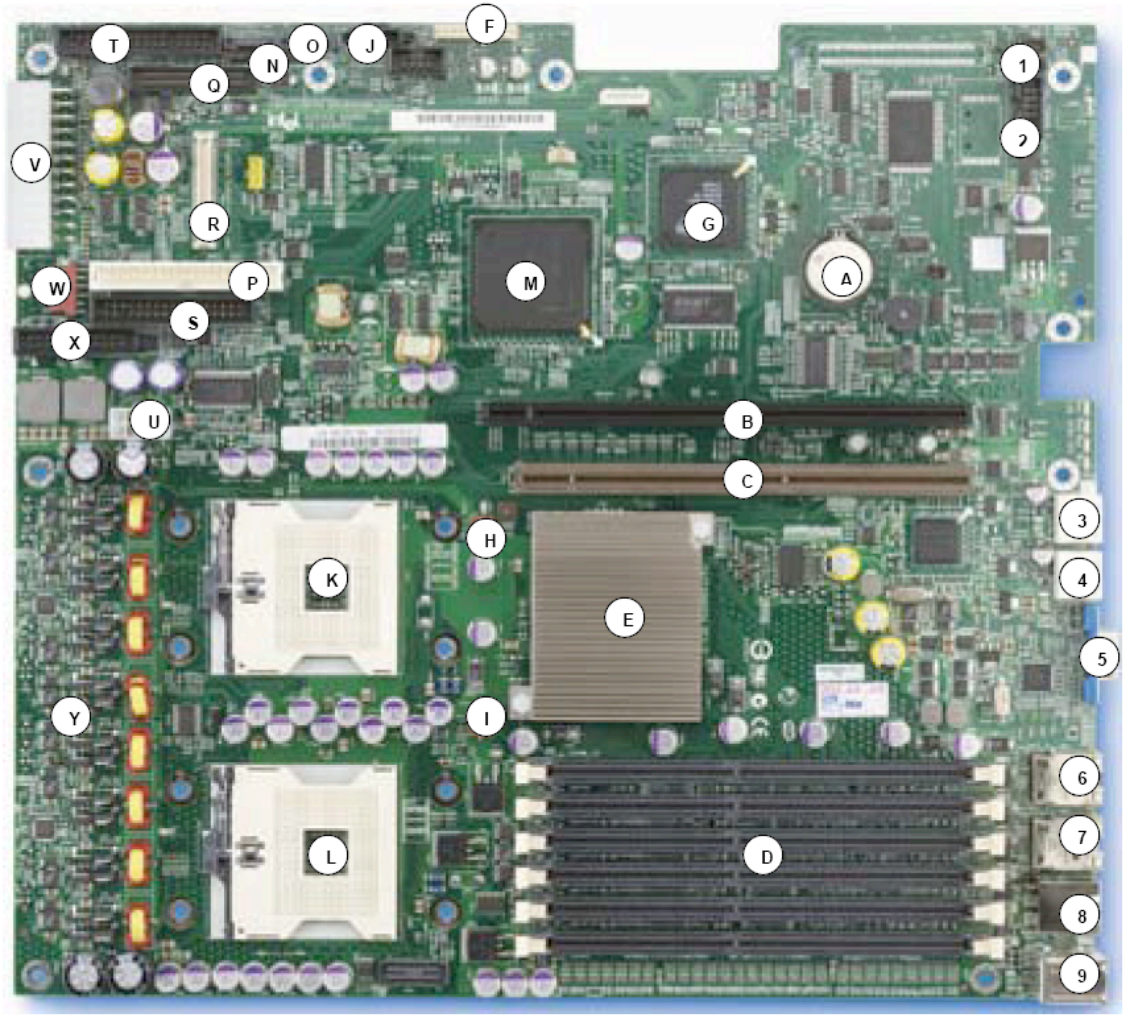

Ref #	Description	Ref #	Description
1	(J1A1) 2-pin Chassis Intrusion Header (J1A2) 2-pin Hard Drive Act LED Header (J1A4) Rolling BIOS Jumper	I	CPU #1 Fan Header
2	10-pin DH10 Serial A Header	J	5-pin Power Sense Header
3	USB Port 2	K	CPU #2 Socket
4	USB Port 1	L	CPU #1 Socket
5	Video Connector	M	6300ESB 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	RJ-45 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 do nó de computação original SE7320VP2 (c0-*)

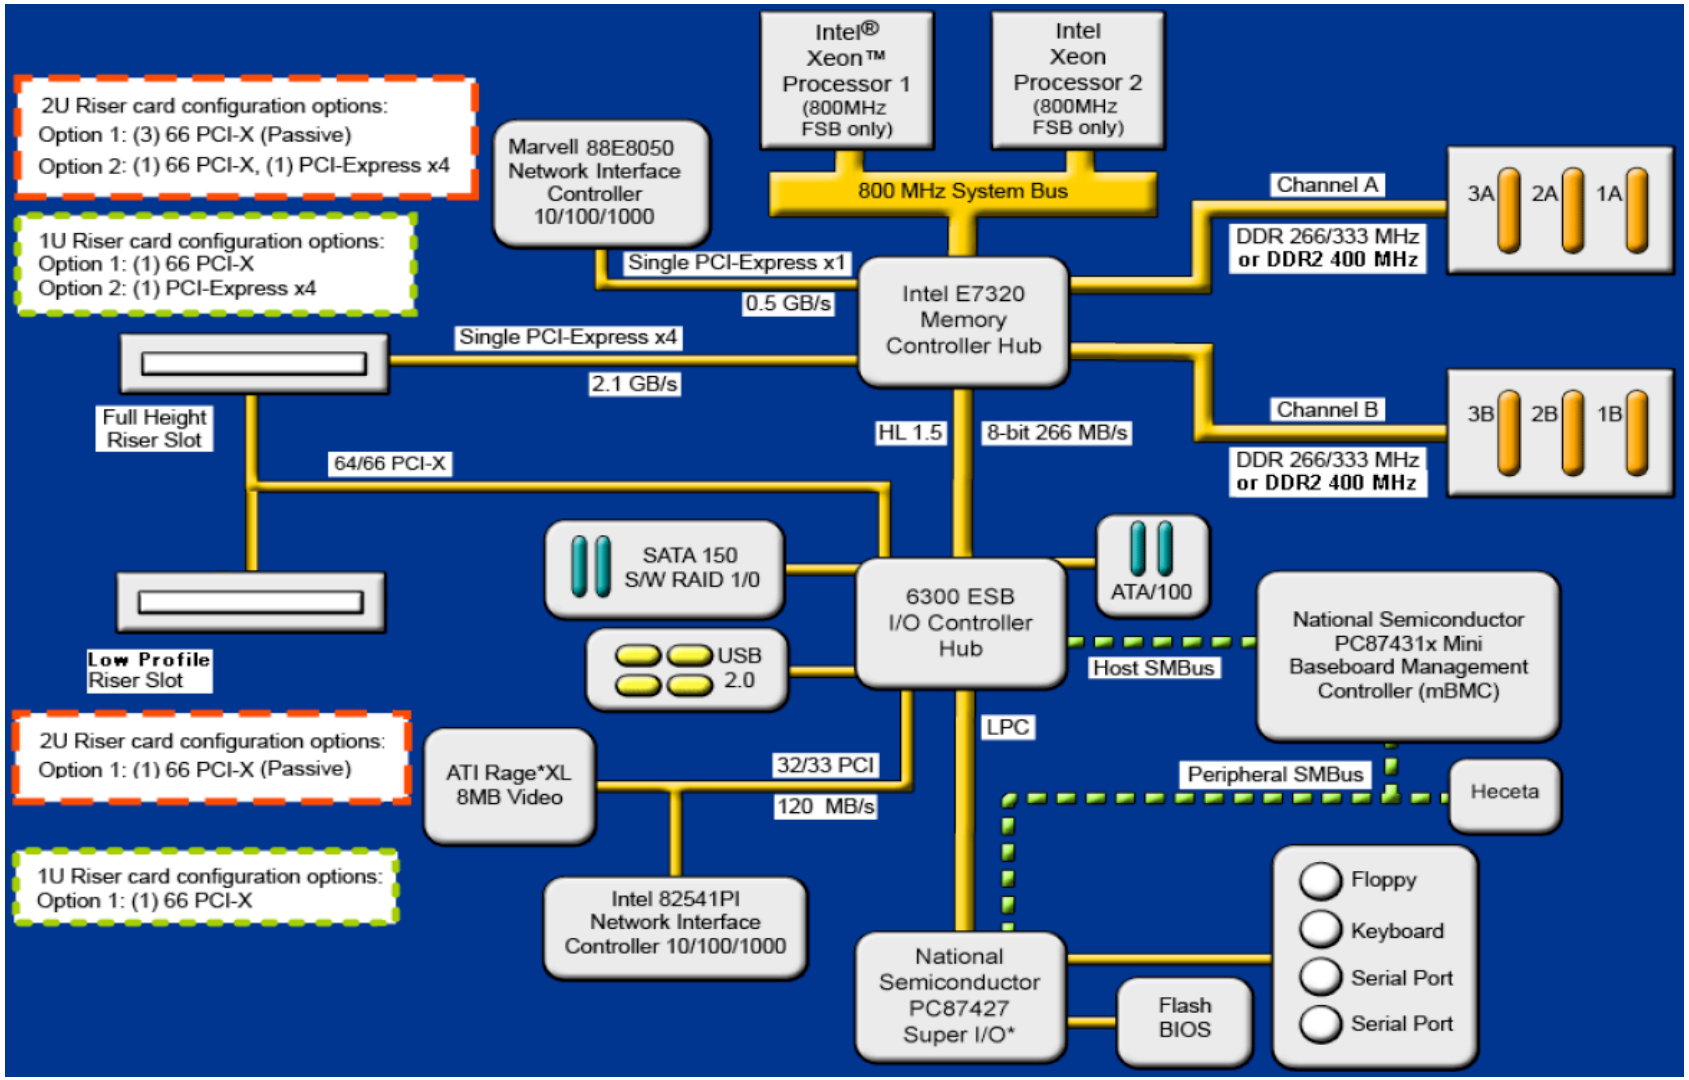

Diagrama do 1º nó de computação dual core SE7520JR2 (c1-1)

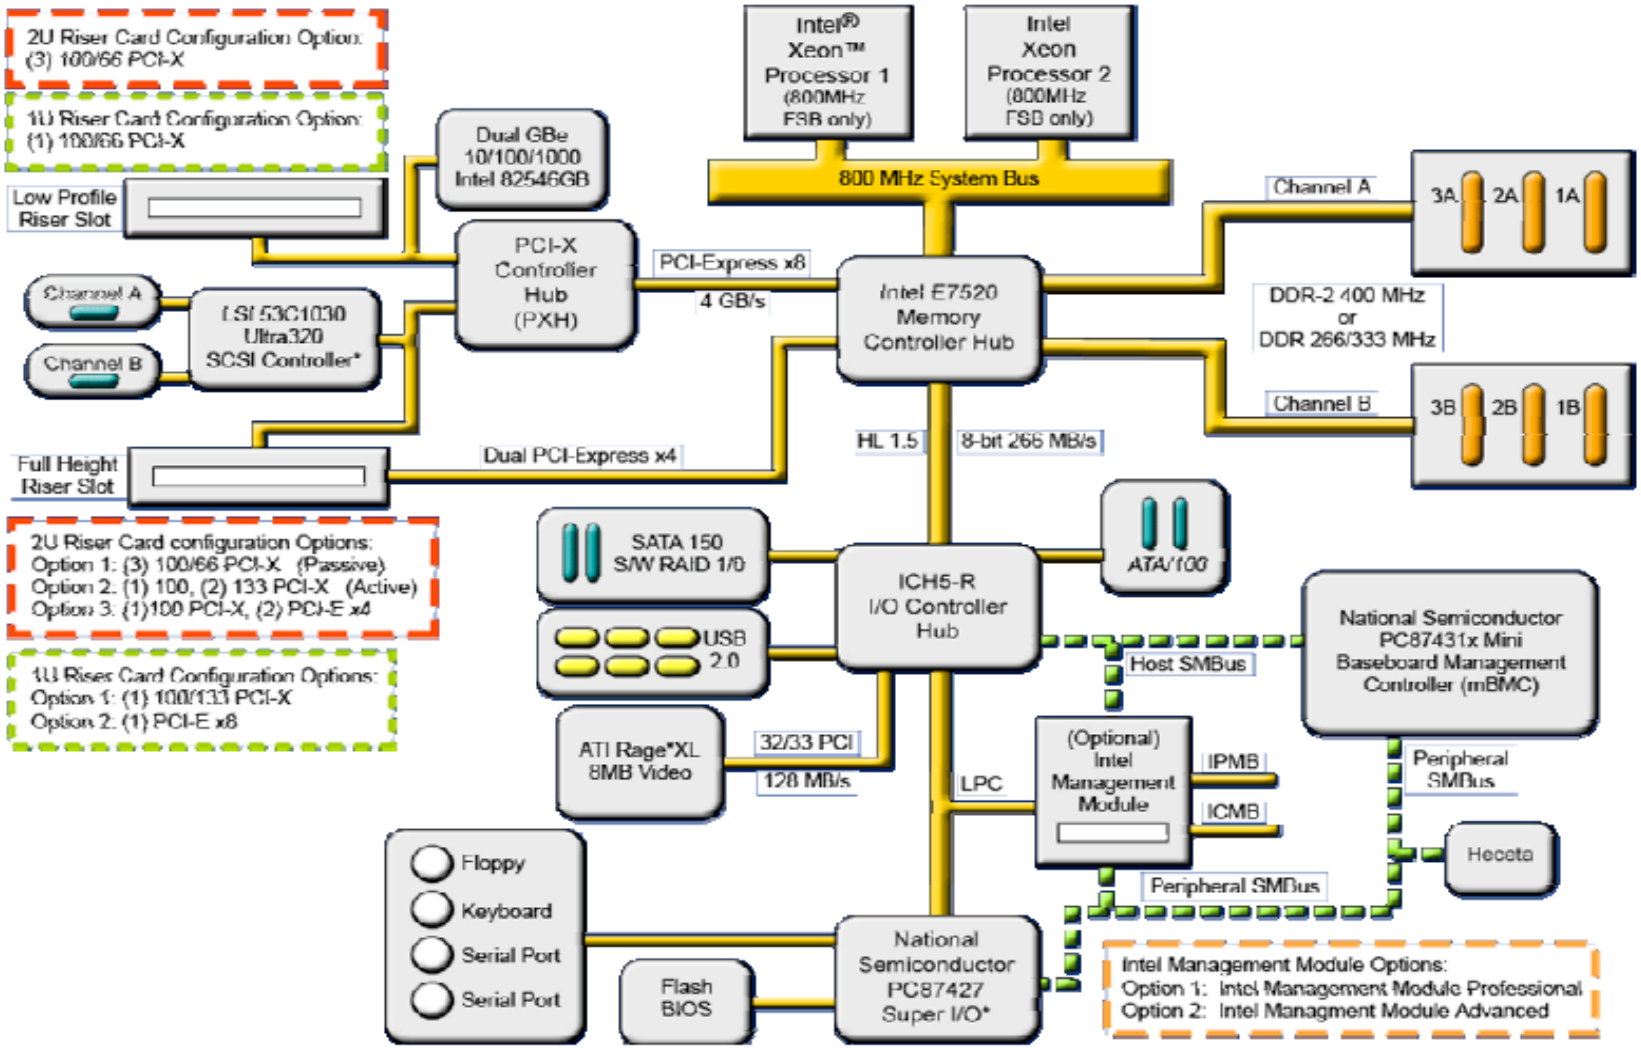

Diagrama dos nós com GPU S5000PAL (c3-* e c6-*)

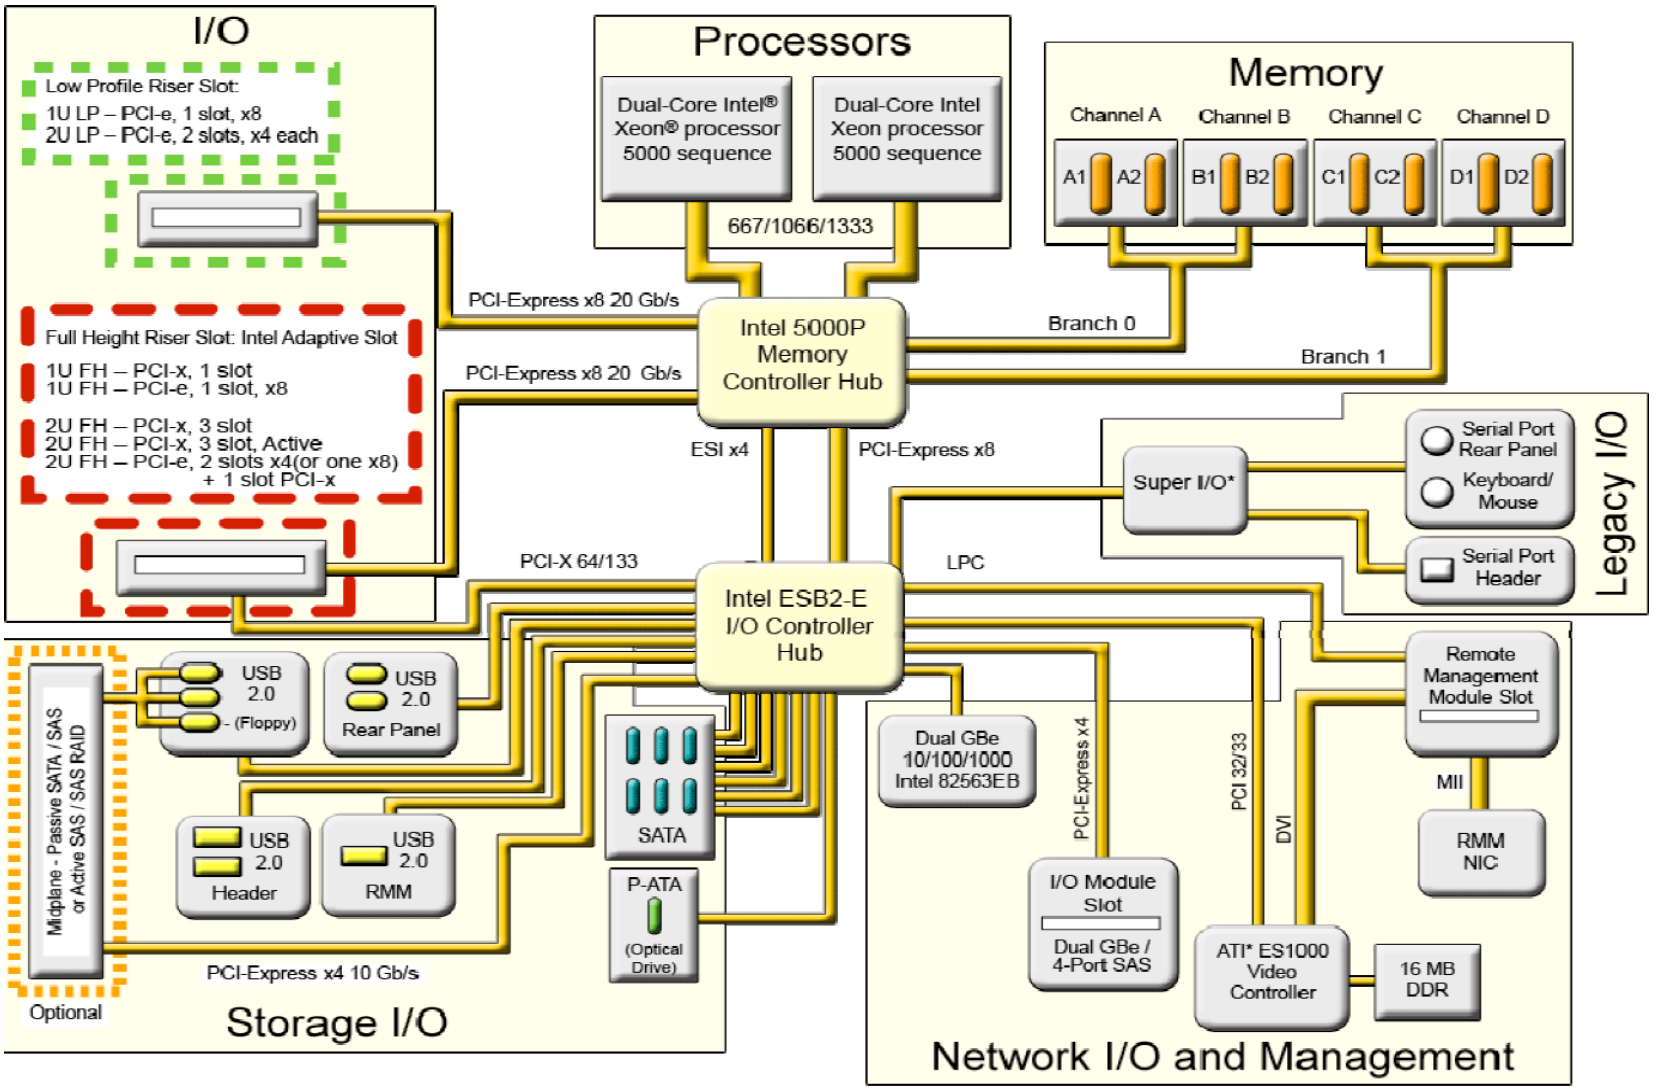

Nó de computação (2009)

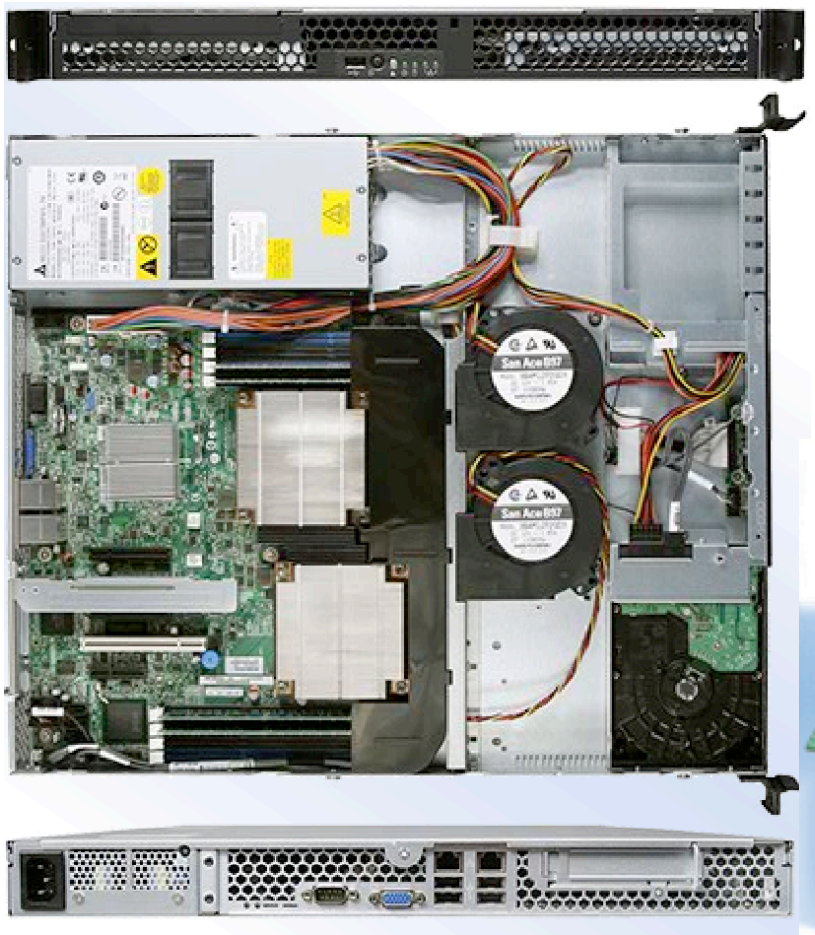

Diagrama funcional da placa mãe Intel S5500BC (c4-*)

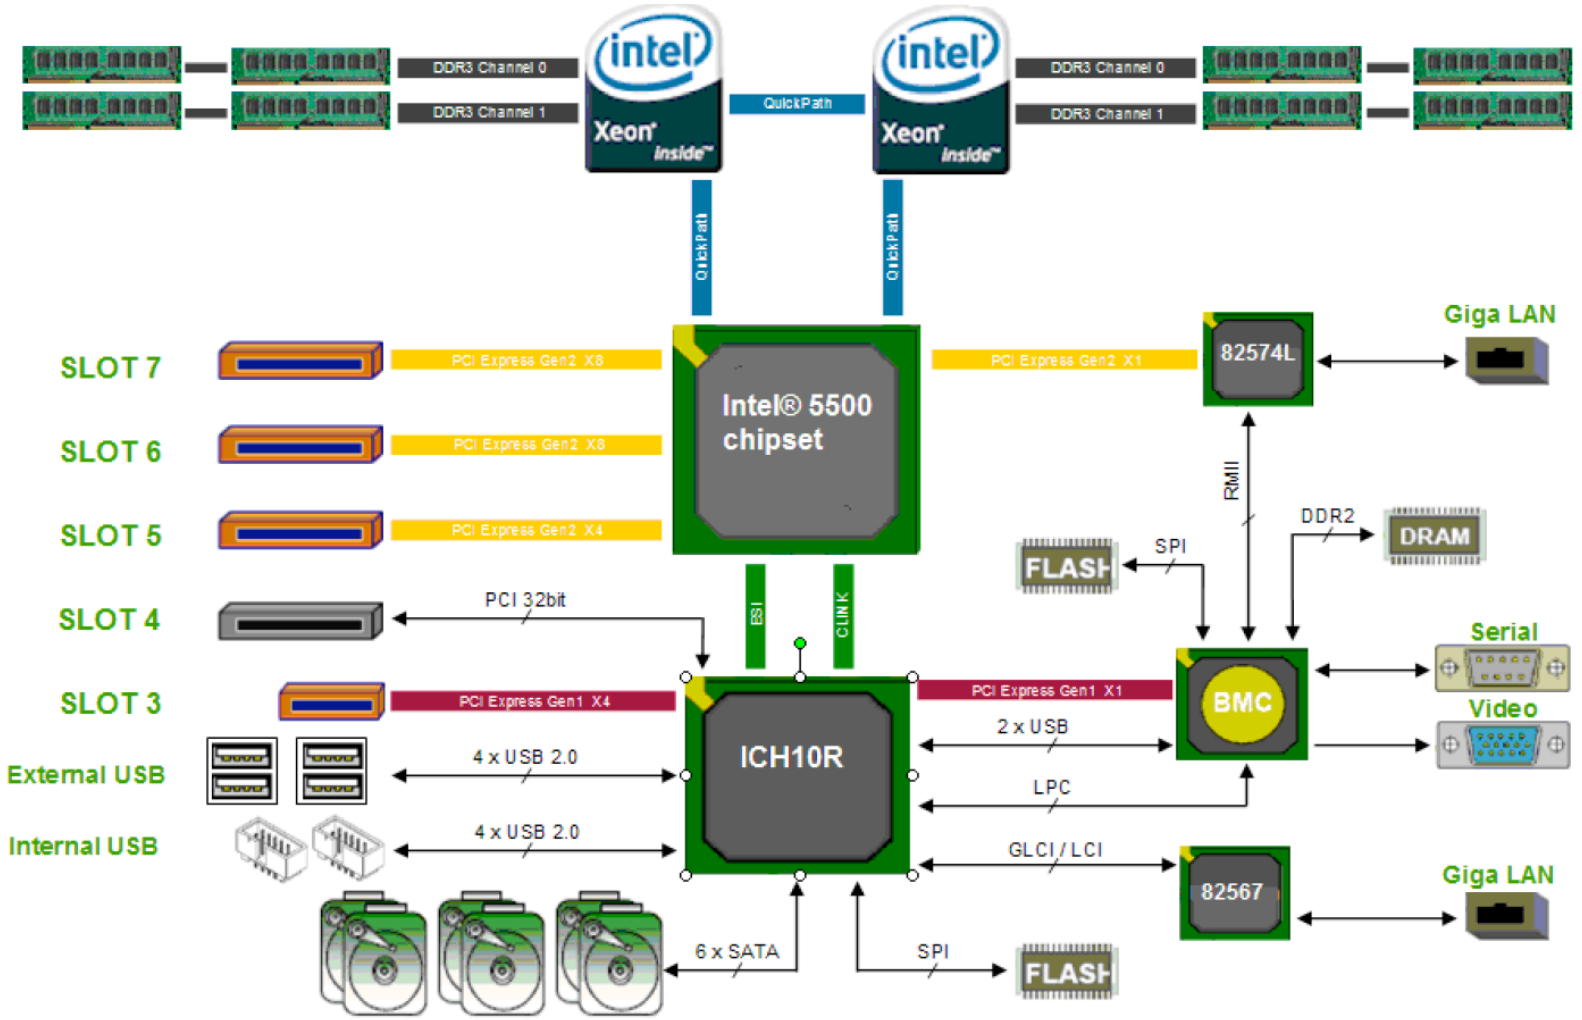

Vista aérea da motherboard S5500BC

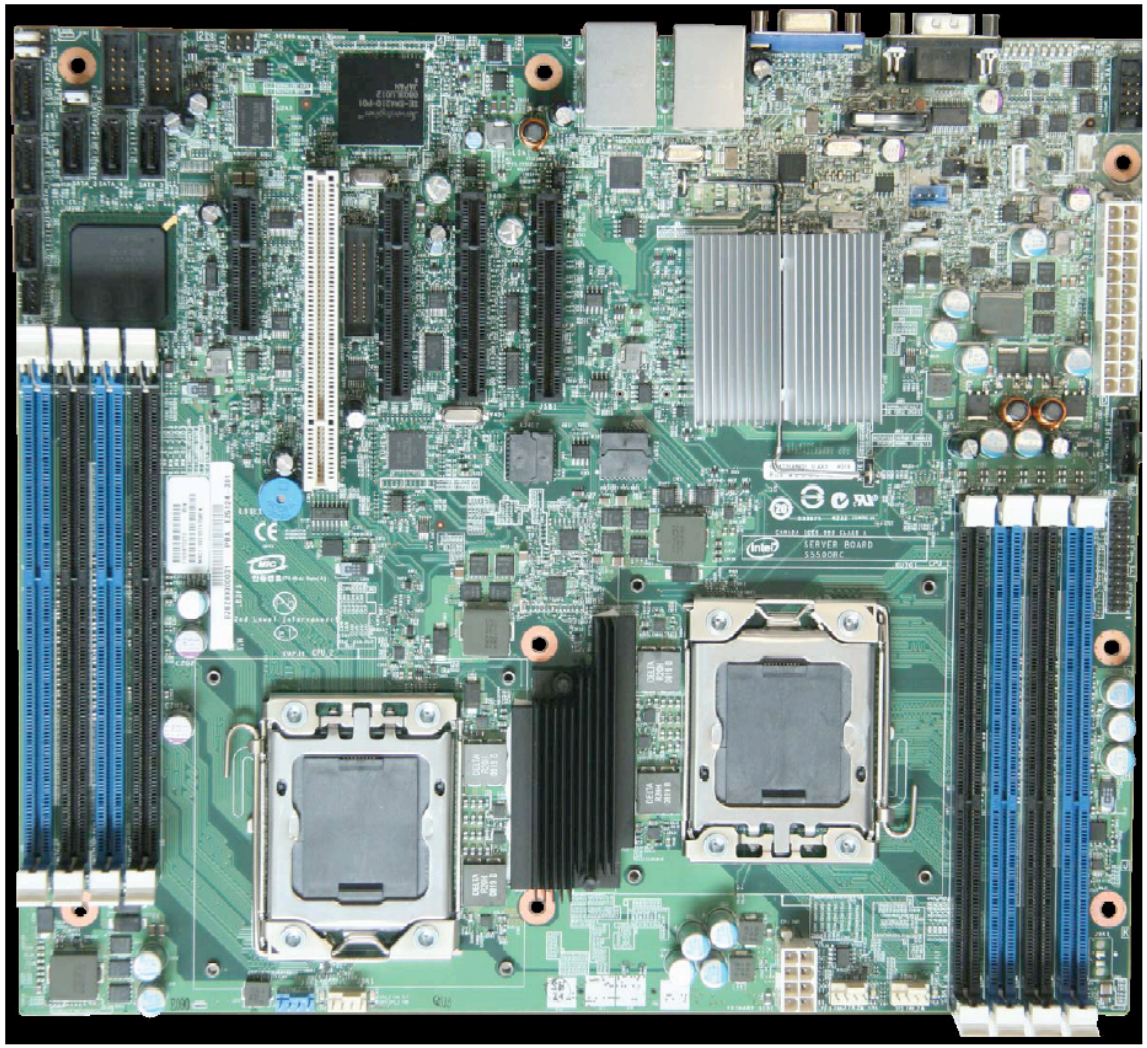

Cluster SeARCH

Evolução da arquitectura lógica

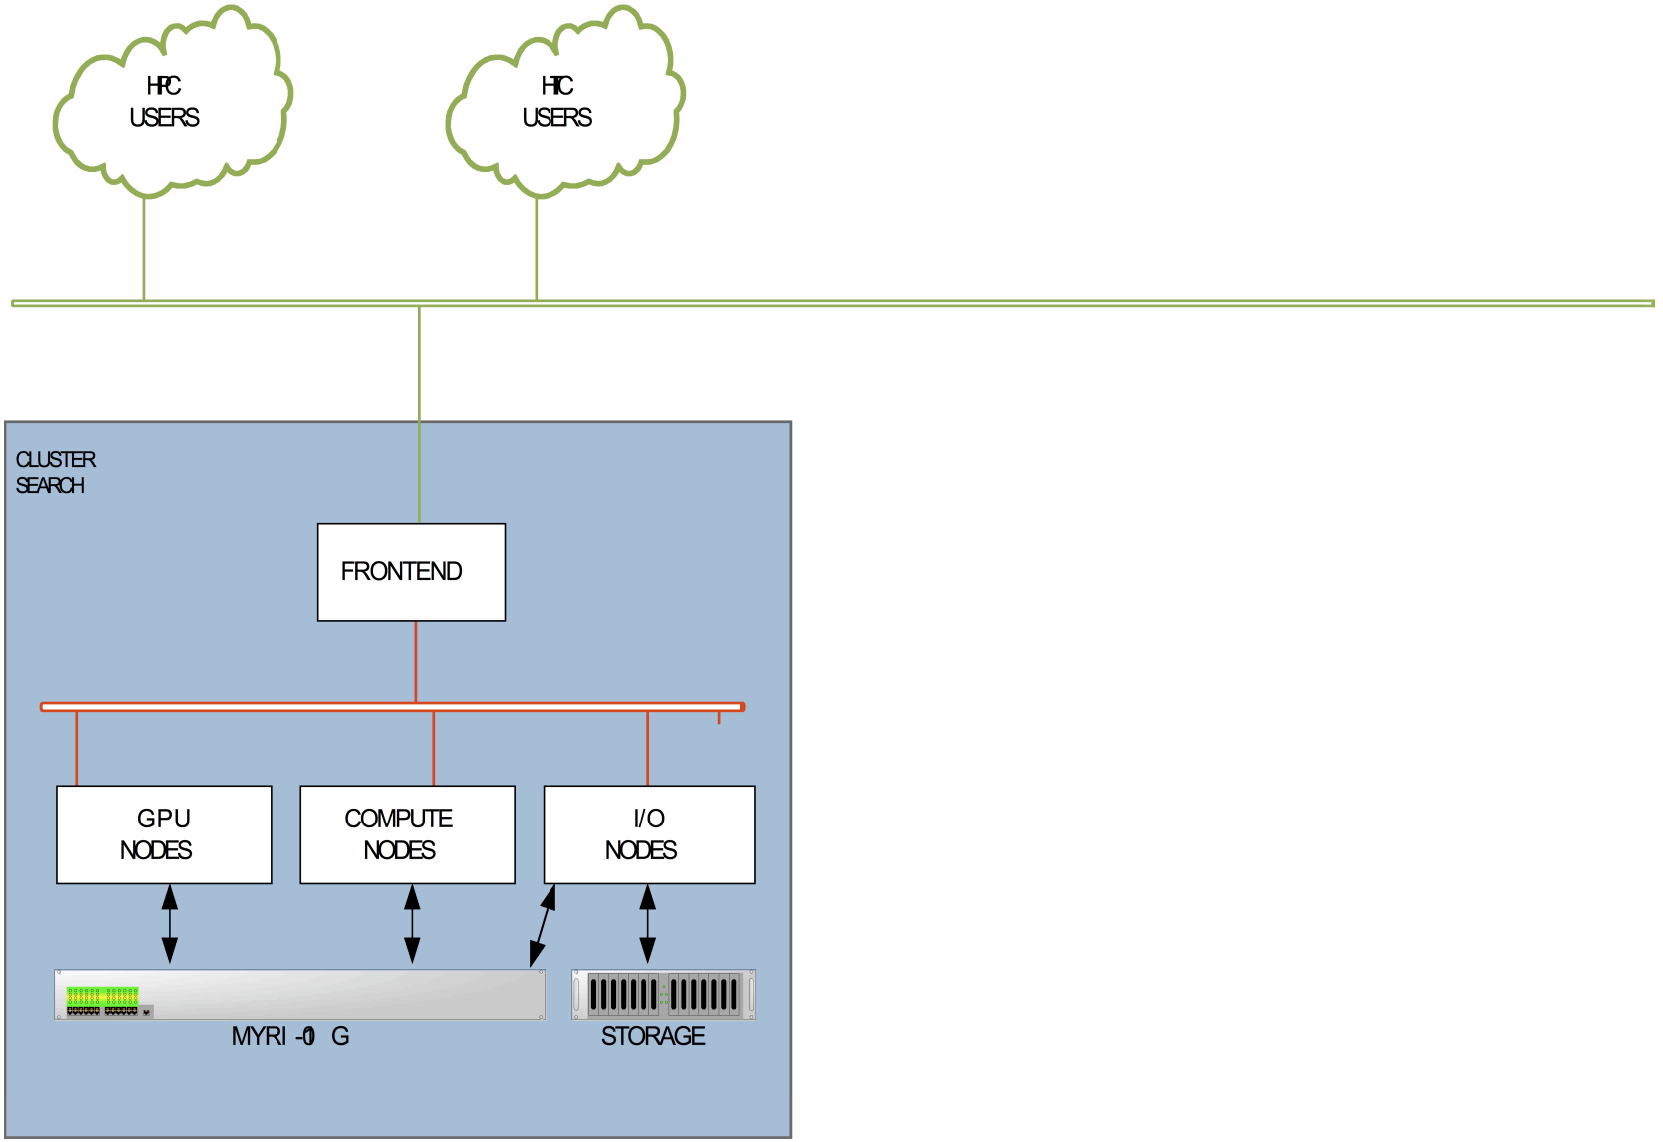
HPC
USERS
HPC
USERS
CLUSTER
SEARCH
FRONTEND
GPU
NODES
COMPUTE
NODES
I/O
NODES
MYRI -0 G
STORAGE

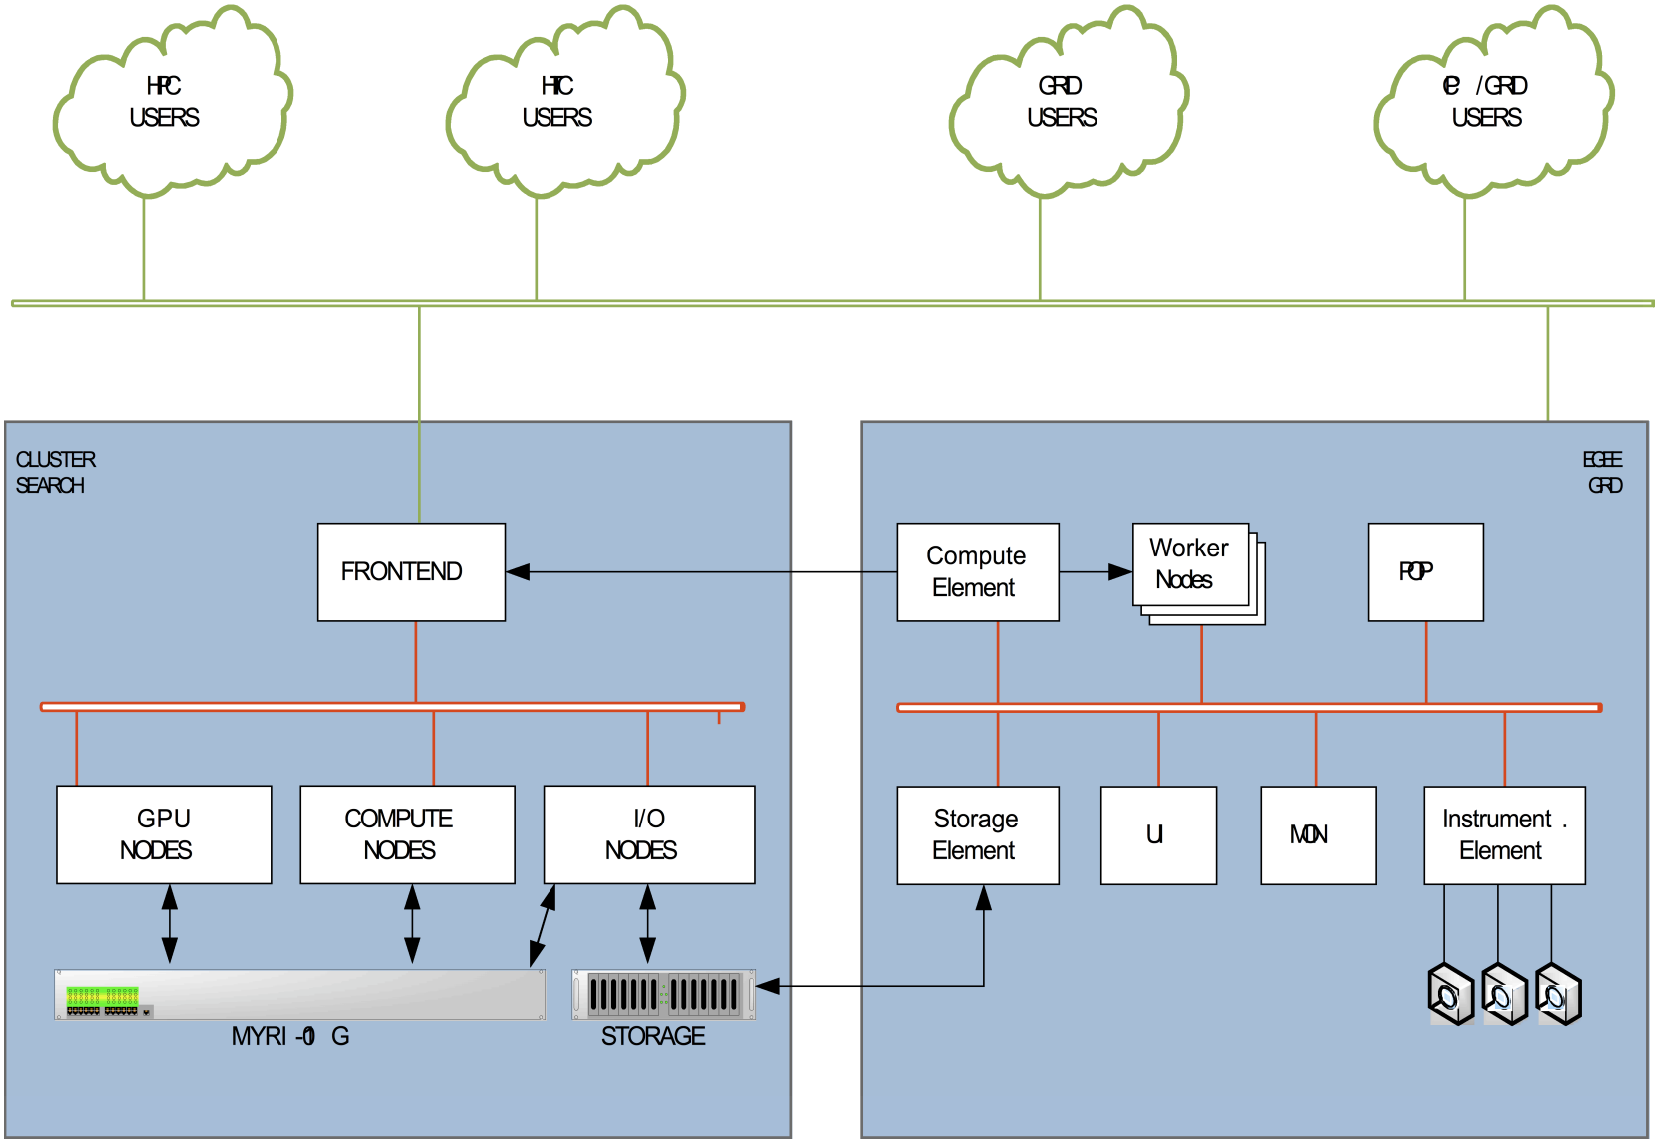
HPC
USERS
HTC
USERS
GRID
USERS
e /GRID
USERS
CLUSTER
SEARCH
FRONTEND
GPU
NODES
COMPUTE
NODES
I/O
NODES
MYRI -0 G
STORAGE
EGEE
GRID
Compute
Element
Worker
Nodes
ROP
Storage
Element
UI
MON
Instrument .
Element

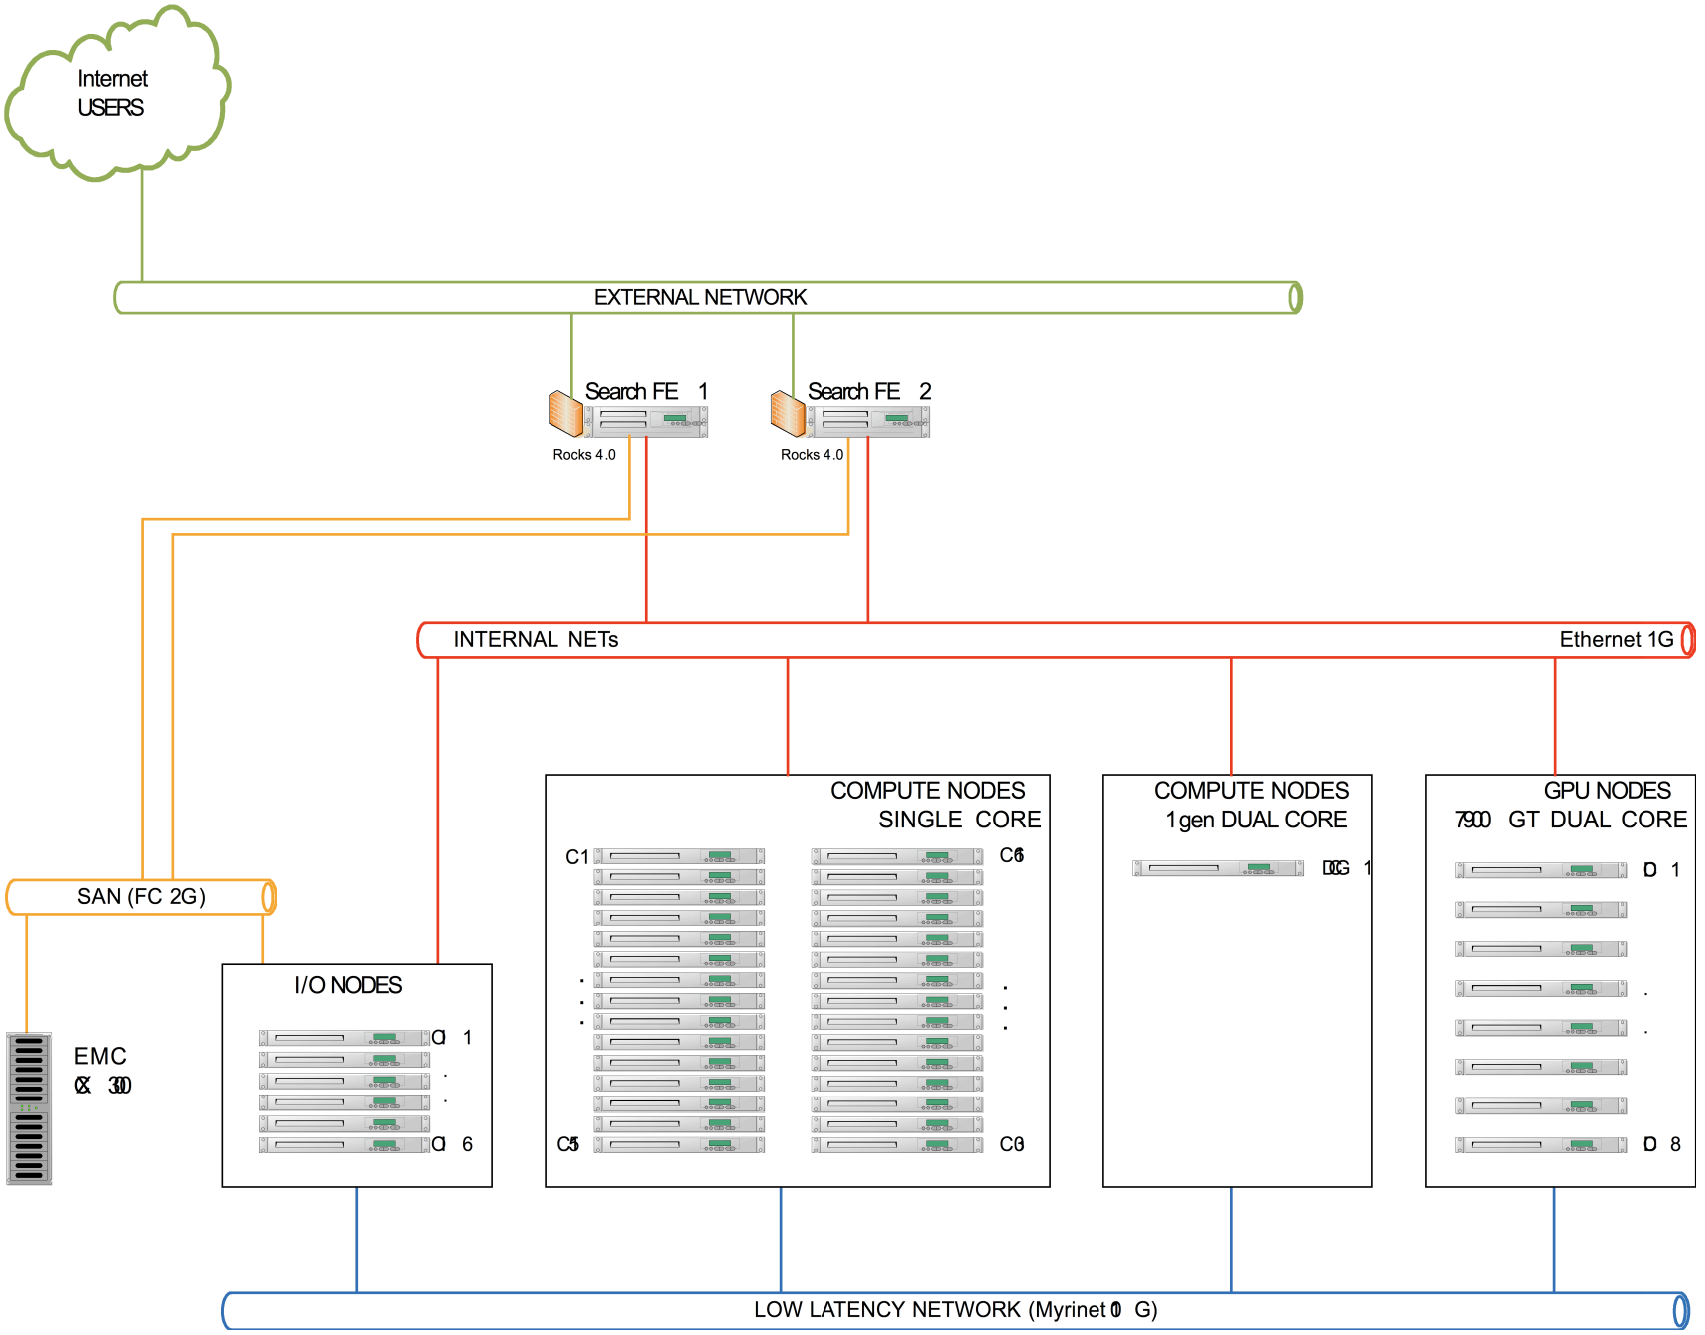
Internet
USERS
EXTERNAL NETWORK
Search FE 1
Rocks 4.0
Search FE 2
Rocks 4.0
INTERNAL NETs
Ethernet 1G
SAN (FC 2G)
EMC
I/O NODES
COMPUTE NODES
SINGLE CORE
COMPUTE NODES
1gen DUAL CORE
GPU NODES
GT DUAL CORE
LOW LATENCY NETWORK (Myrinet G)

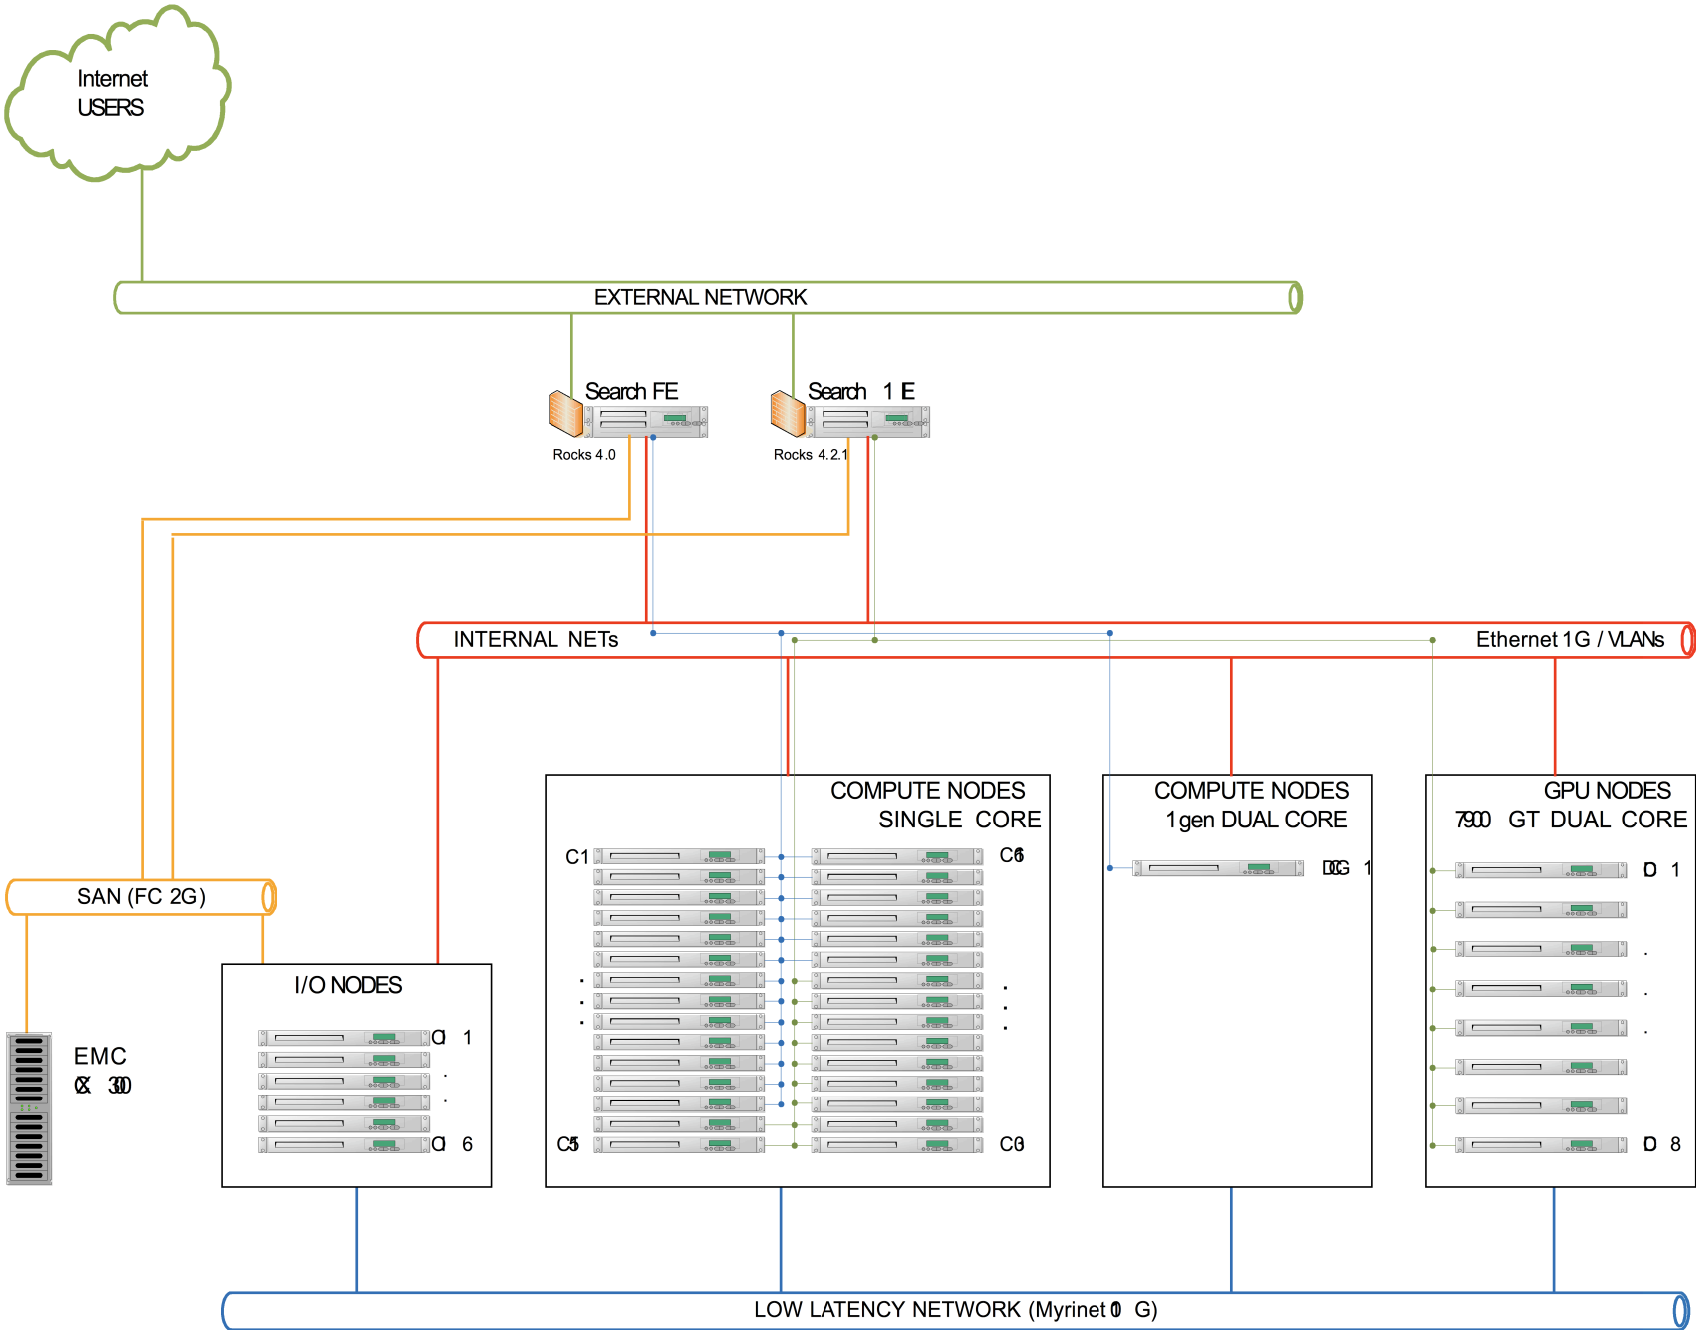
Internet
USERS
EXTERNAL NETWORK
Search FE
Rocks 4.0
Search 1 E
Rocks 4.2.1
INTERNAL NETs
Ethernet 1G / VLANs
SAN (FC 2G)
EMC
I/O NODES
COMPUTE NODES
SINGLE CORE
COMPUTE NODES
1gen DUAL CORE
GPU NODES
GT DUAL CORE
LOW LATENCY NETWORK (Myrinet 0 G)

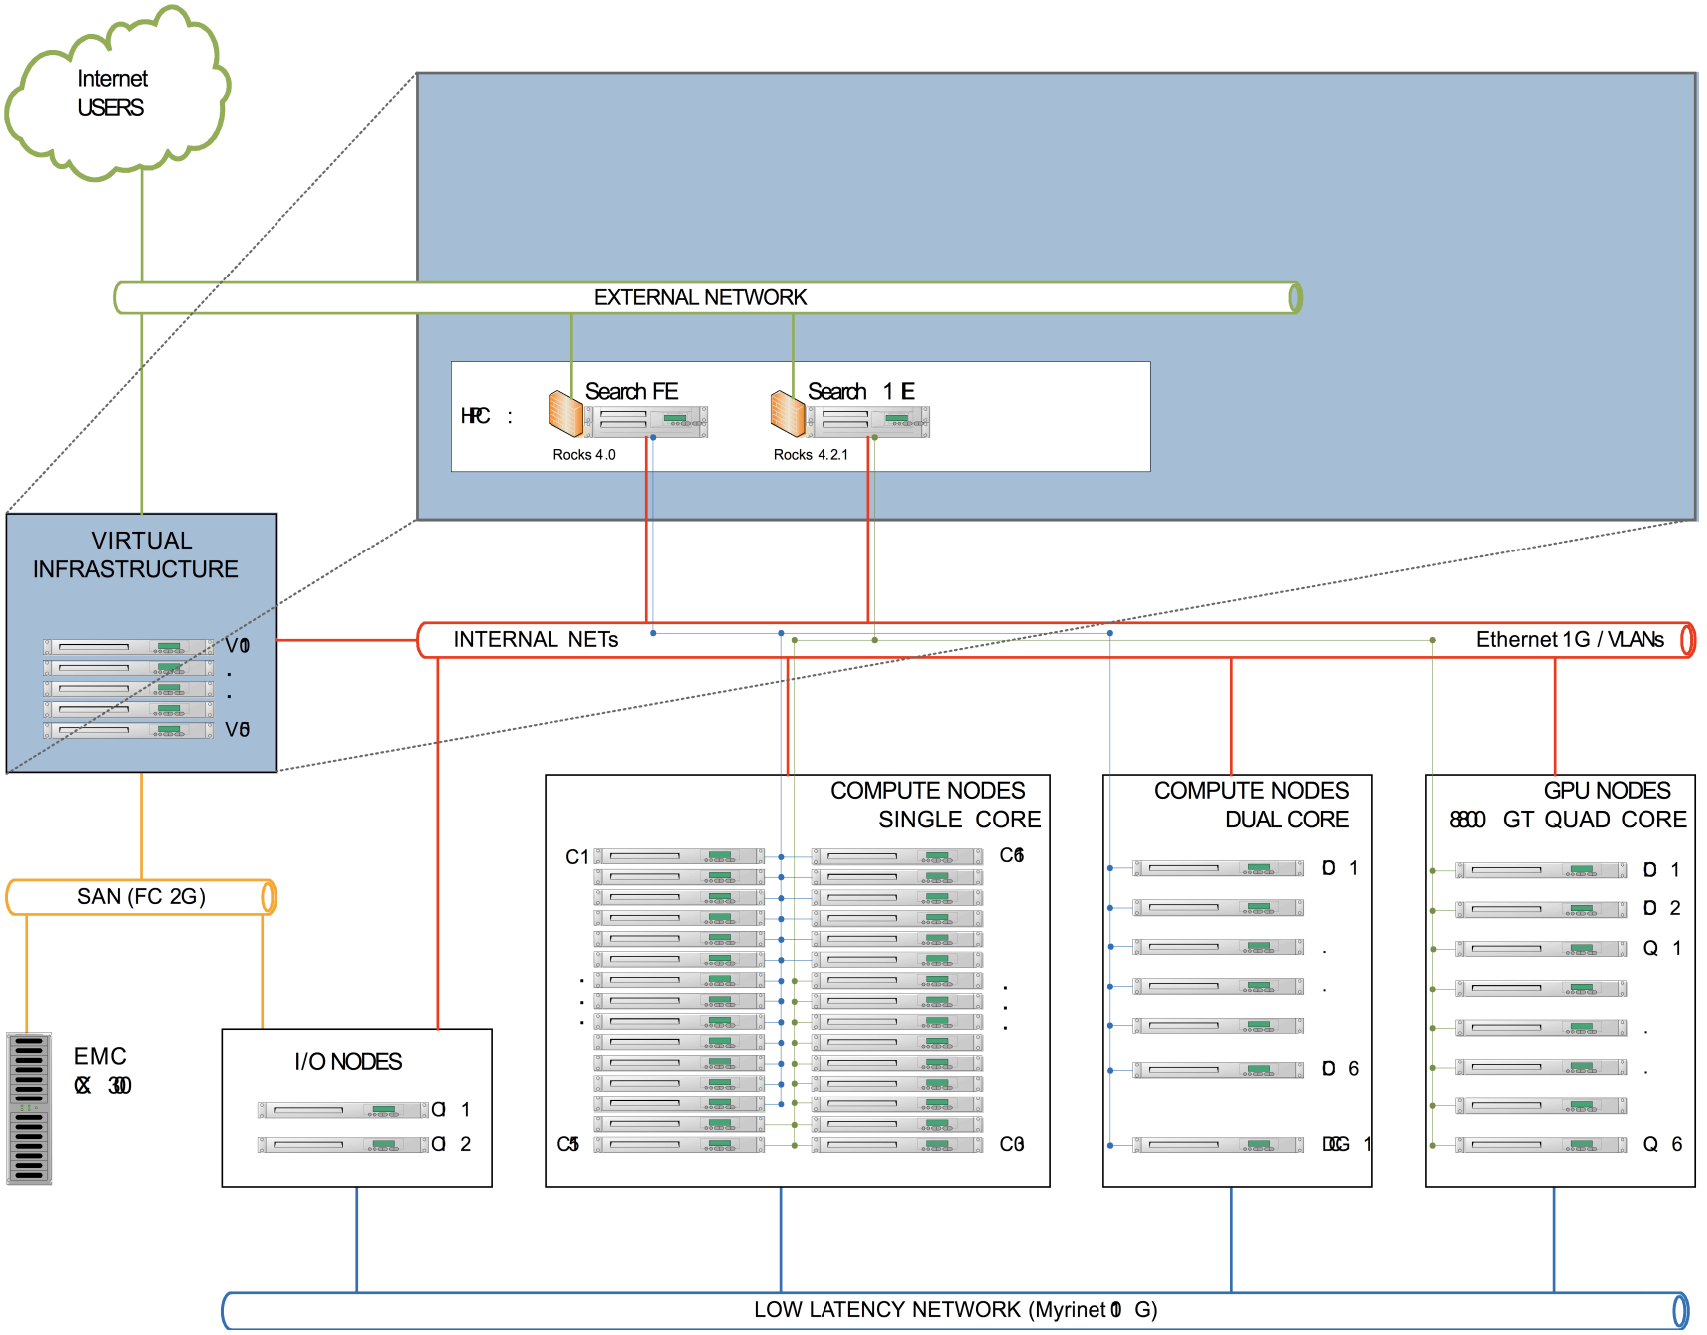
Internet
USERS
EXTERNAL NETWORK
Search FE
Rocks 4.0
Rocks 4.2.1
VIRTUAL
INFRASTRUCTURE
INTERNAL NETs
Ethernet 1G / VLANs
SAN (FC 2G)
EMC
I/O NODES
COMPUTE NODES
SINGLE CORE
COMPUTE NODES
DUAL CORE
GPU NODES
GT QUAD CORE
LOW LATENCY NETWORK (Myrinet 10 G)

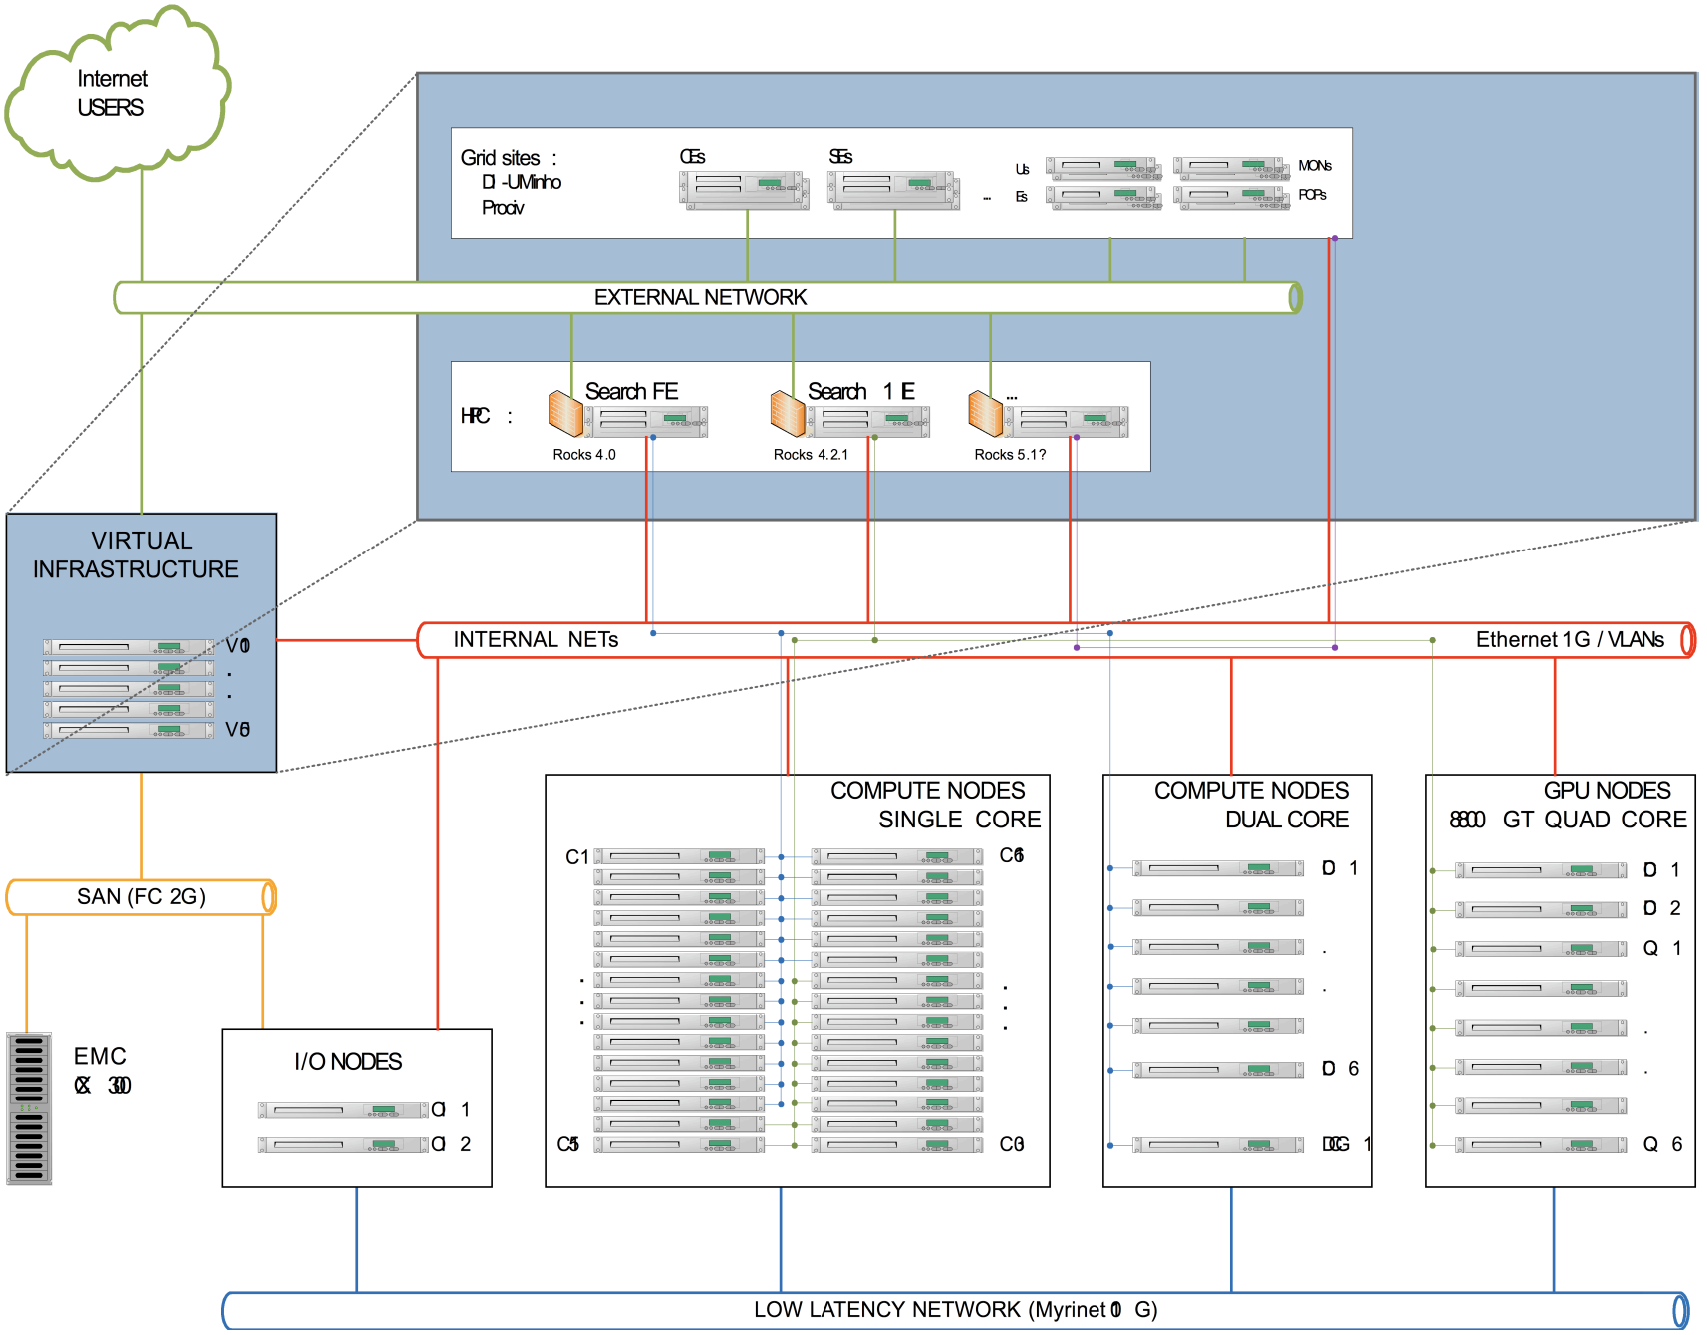
Internet
USERS
Grid sites :
CEs
SEs
MONs
PCPs
EXTERNAL NETWORK
HPC :
Search FE
Search 1 E
...
Rocks 4.0
Rocks 4.2.1
Rocks 5.1?
VIRTUAL
INFRASTRUCTURE
INTERNAL NETs
Ethernet 1G / VLANs
SAN (FC 2G)
EMC
I/O NODES
COMPUTE NODES
SINGLE CORE
COMPUTE NODES
DUAL CORE
GPU NODES
GT QUAD CORE
LOW LATENCY NETWORK (Myrinet G)

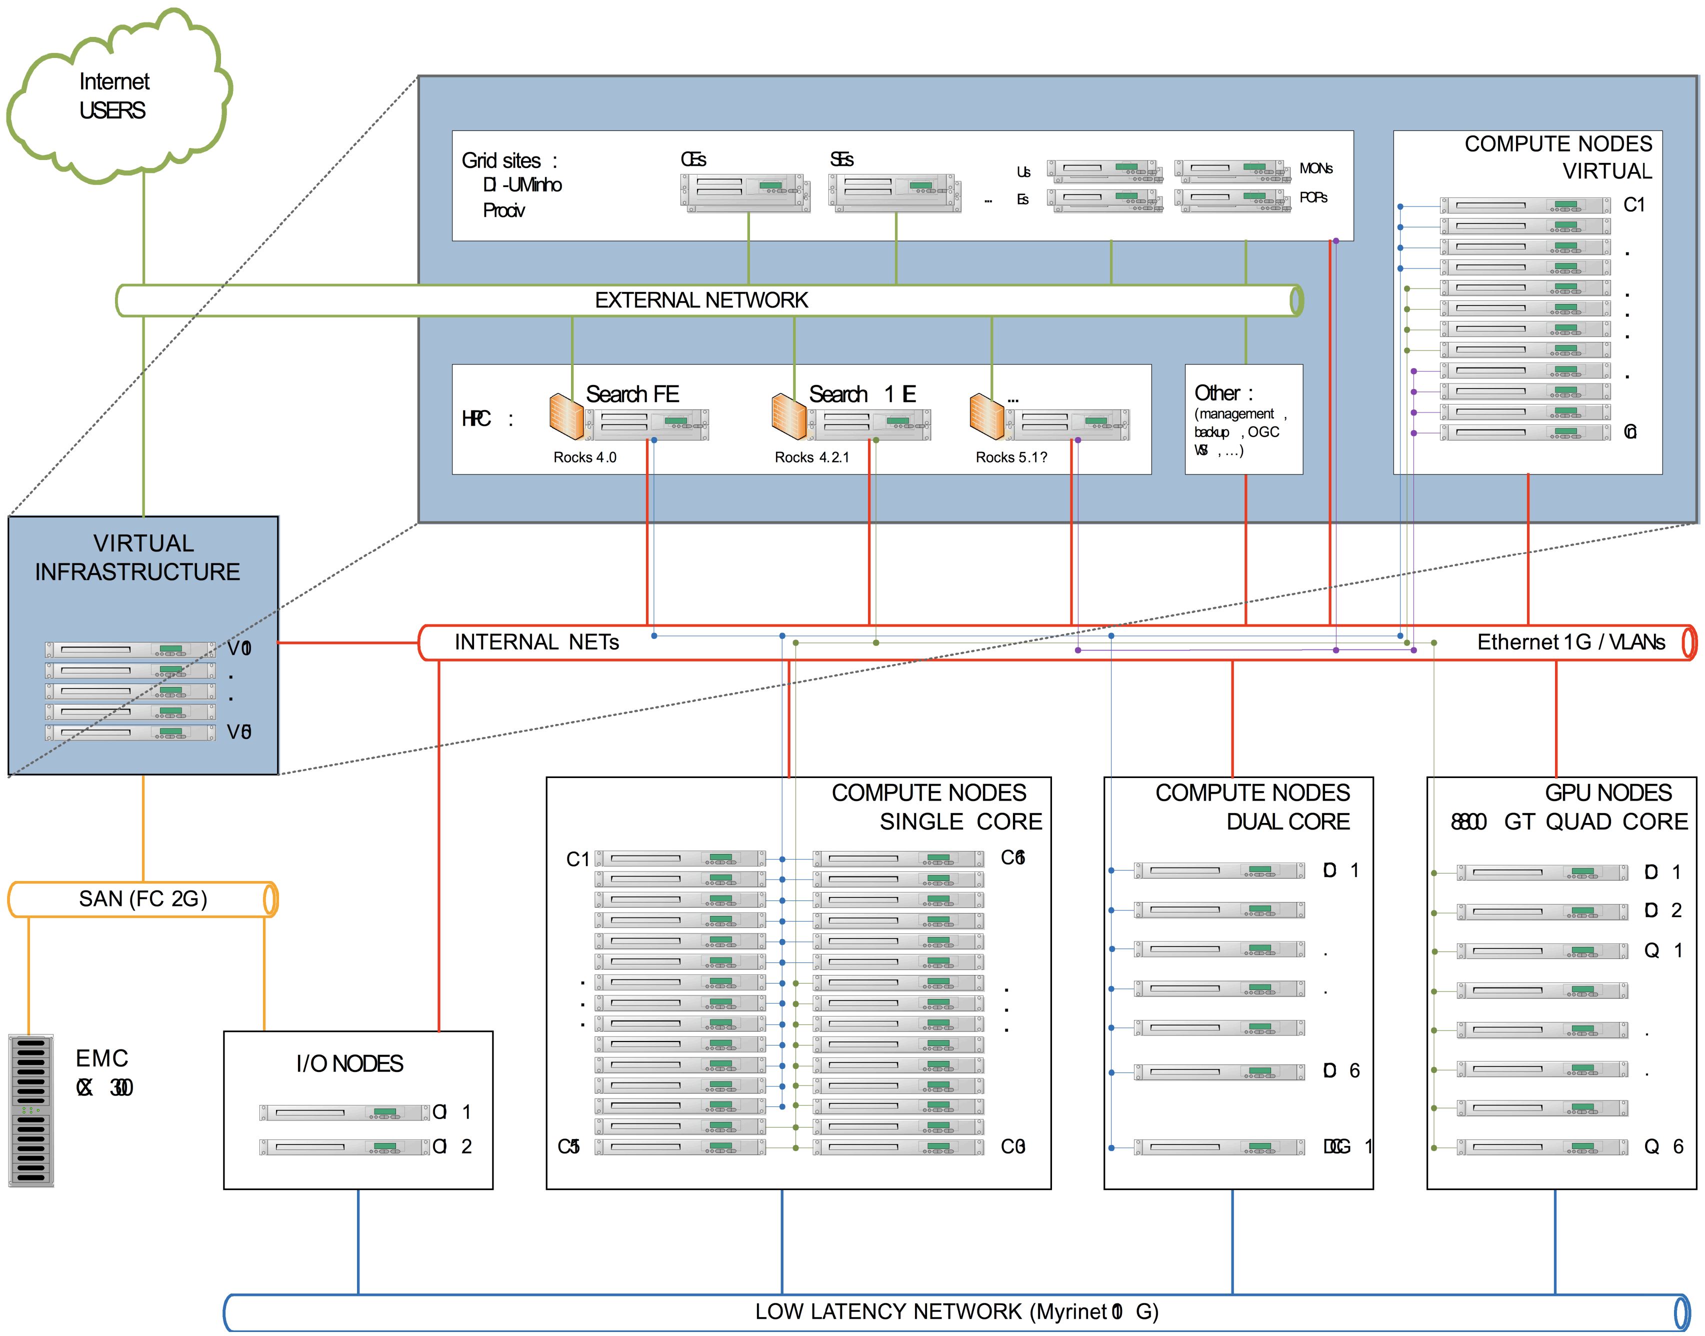

Internet
USERS
Grid sites :
Prociv
EXTERNAL NETWORK
HPC :
Search FE
Rocks 4.0
Rocks 4.2.1
Rocks 5.1?
Other :
(management ,
backup , OGC
, ...)
COMPUTE NODES
VIRTUAL
C1
VIRTUAL
INFRASTRUCTURE
INTERNAL NETs
Ethernet 1G / VLANs
COMPUTE NODES
SINGLE CORE
C1
COMPUTE NODES
DUAL CORE
GPU NODES
GT QUAD CORE
SAN (FC 2G)
EMC
I/O NODES
LOW LATENCY NETWORK (Myrinet
G)

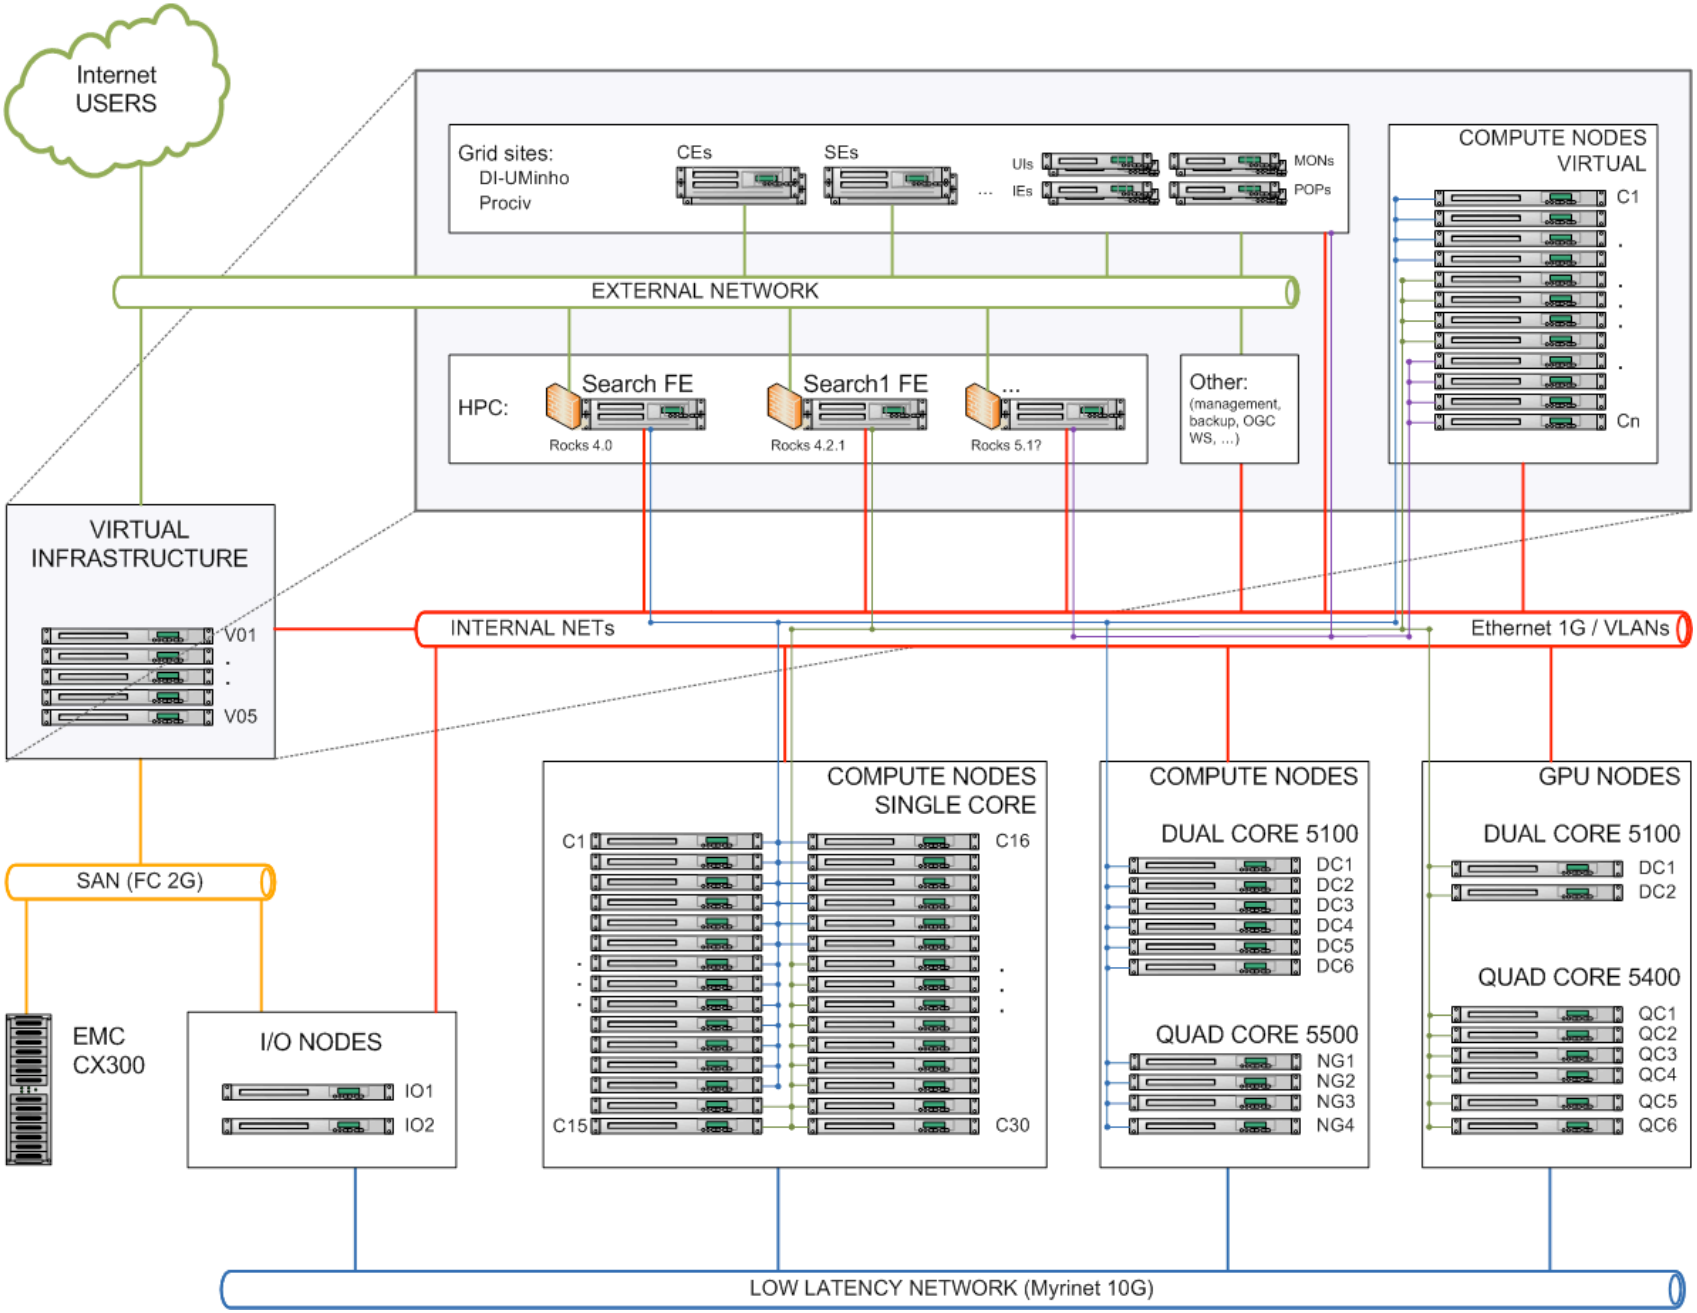
Internet
USERS
Grid sites:
DI-UMinho
Prociv
CEs
SEs
UIs
IEs
MONs
POPs
COMPUTE NODES
VIRTUAL
C1
Cn
EXTERNAL NETWORK
HPC:
Search FE
Search1 FE
...
Rocks 4.0
Rocks 4.2.1
Rocks 5.1?
Other:
(management,
backup, OGC
WS, ...)
VIRTUAL
INFRASTRUCTURE
V01
V05
INTERNAL NETs
Ethernet 1G / VLANs
COMPUTE NODES
SINGLE CORE
C1
C16
C15
C30
COMPUTE NODES
DUAL CORE 5100
DC1
DC2
DC3
DC4
DC5
DC6
QUAD CORE 5500
NG1
NG2
NG3
NG4
GPU NODES
DUAL CORE 5100
DC1
DC2
QUAD CORE 5400
QC1
QC2
QC3
QC4
QC5
QC6
SAN (FC 2G)
EMC
CX300
I/O NODES
IO1
IO2
LOW LATENCY NETWORK (Myrinet 10G)

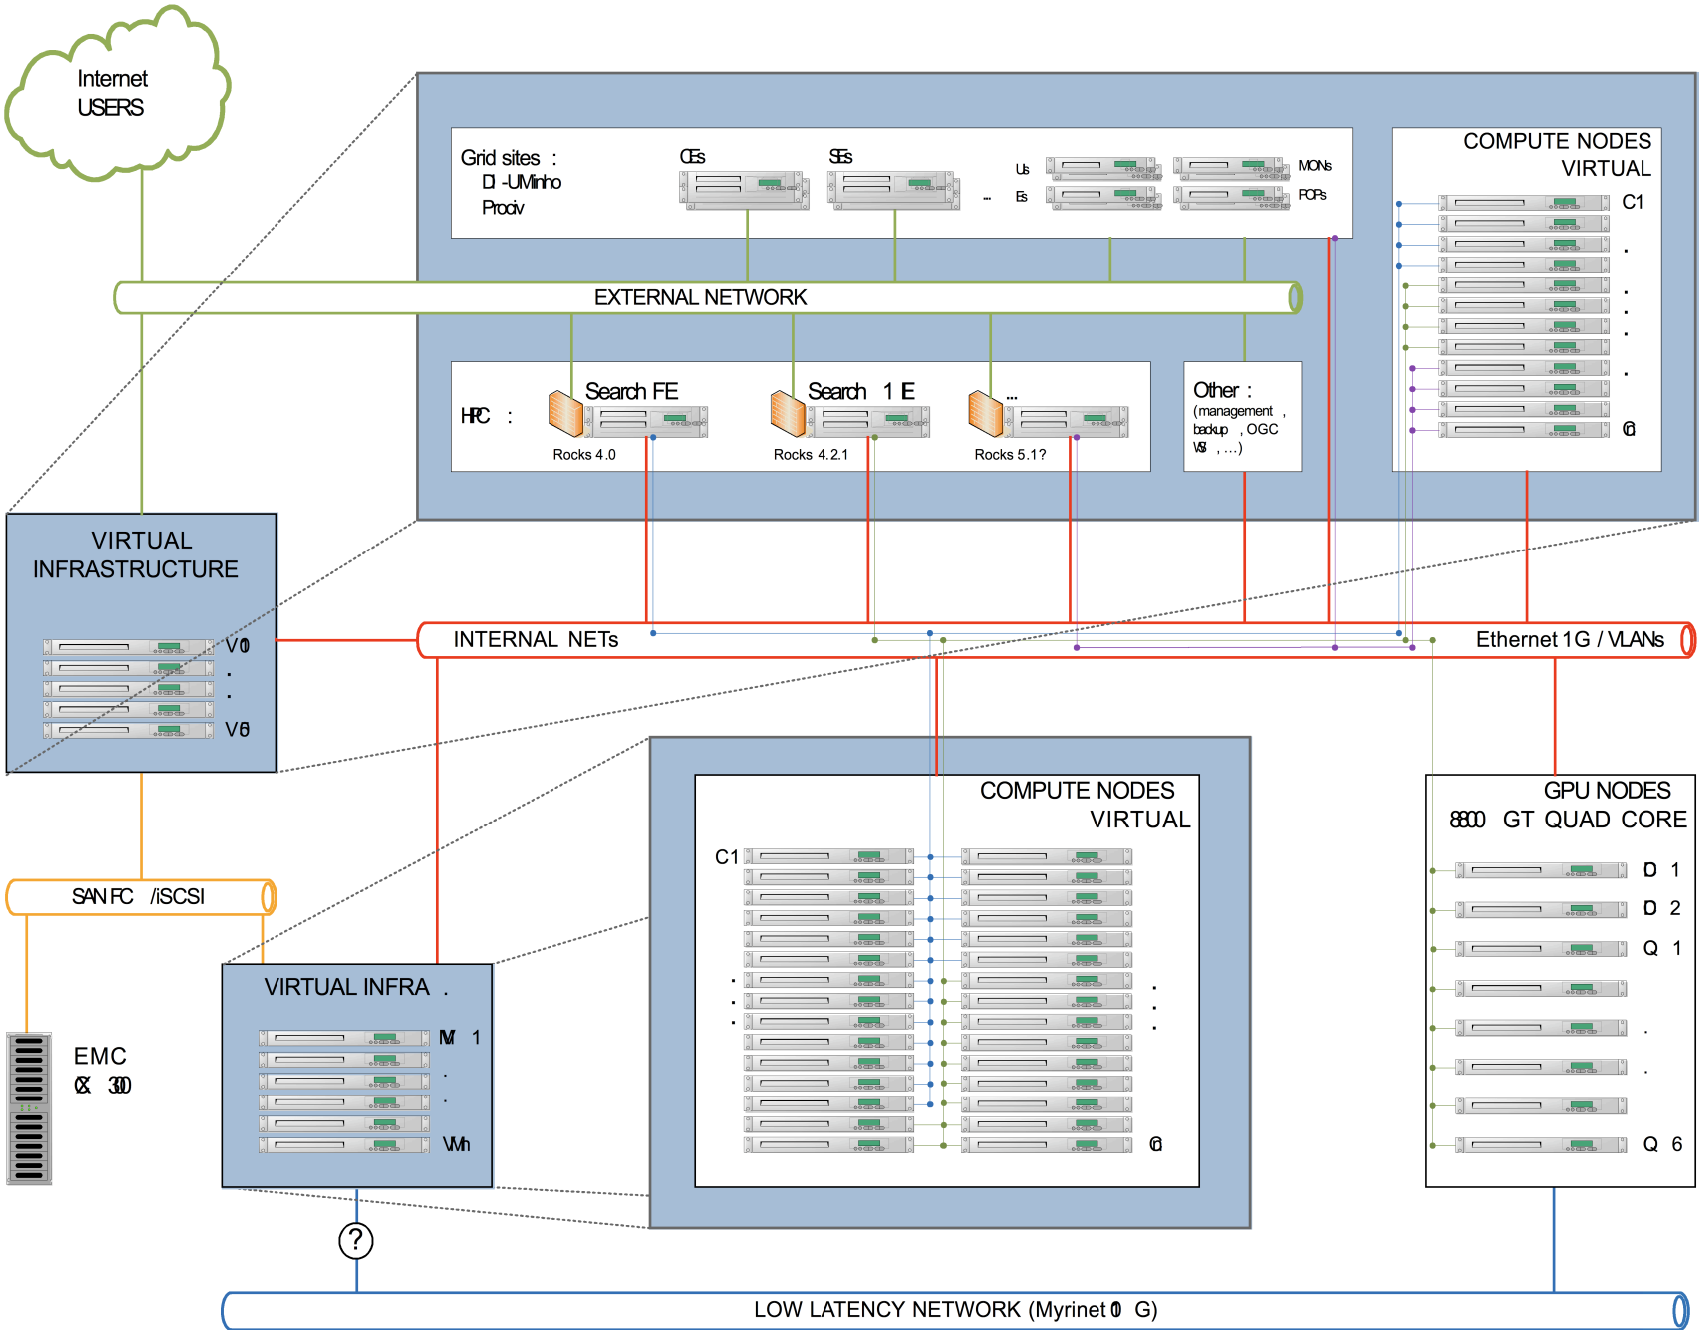
Internet
USERS
Grid sites :
CEs
SEs
MONs
PCPs
COMPUTE NODES
VIRTUAL
C1
EXTERNAL NETWORK
HPC :
Search FE
Search 1 E
...
Rocks 4.0
Rocks 4.2.1
Rocks 5.1?
Other :
(management ,
backup , OGC
VOs , ...)
VIRTUAL
INFRASTRUCTURE
INTERNAL NETs
Ethernet 1G / VLANs
SAN FC /iSCSI
EMC
VIRTUAL INFRA .
COMPUTE NODES
VIRTUAL
C1
GPU NODES
8800 GT QUAD CORE
D 1
D 2
Q 1
Q 6
LOW LATENCY NETWORK (Myrinet 10 G)

Sistemas actuais

Top 500

Junho de 2009

Distribuição por Arquitectura

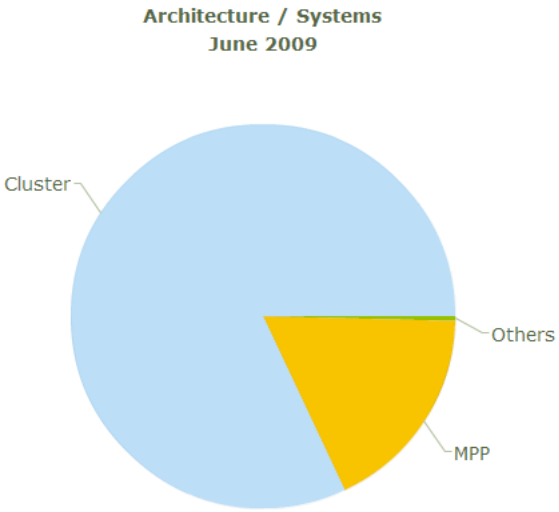

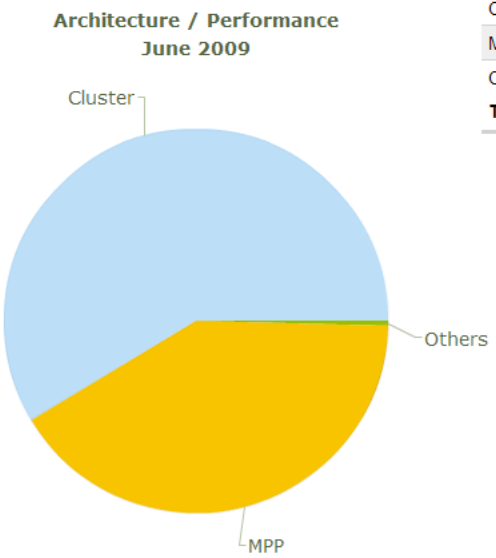

Architecture	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Constellations	2	0.40 %	94970	112947	17648
MPP	88	17.60 %	9255003	11542667	1964865
Cluster	410	82.00 %	13258024	22015991	2122220
Totals	**500**	**100%**	**22607996.30**	**33671604.92**	**4104733**

Distribuição por Tipo de Processador

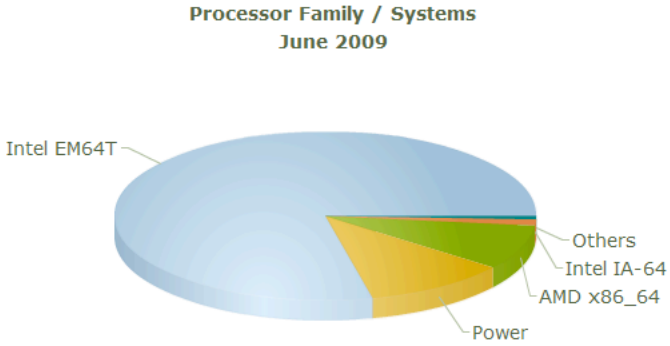

Processor Family	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Power	55	11.00 %	5916480	7451856	1505104
NEC	1	0.20 %	122400	131072	1280
Sparc	1	0.20 %	110600	121282	12032
Intel IA-64	6	1.20 %	289948	343346	54512
Intel EM64T	393	78.60 %	11435904	19291260	1777732
AMD x86_64	43	8.60 %	4710704	6248309	745881
Others	1	0.20 %	21960	84480	8192
Totals	**500**	**100%**	**22607996.30**	**33671604.92**	**4104733**

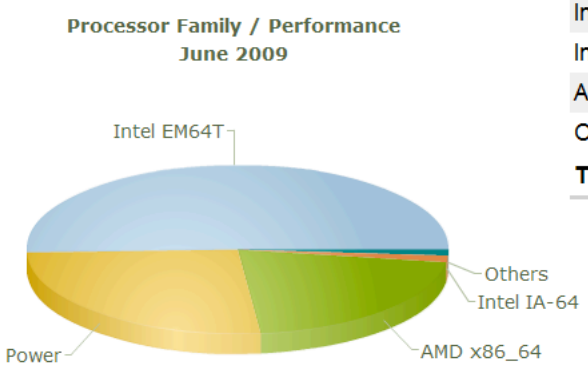

Distribuição por Modelo de Processador

Processor Generation / Systems
June 2009

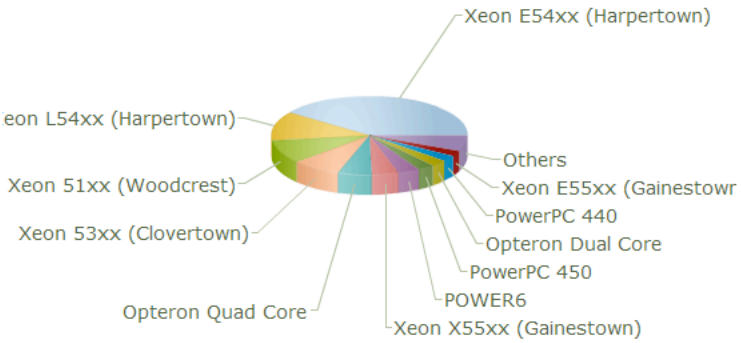

Processor Generation / Performance
June 2009

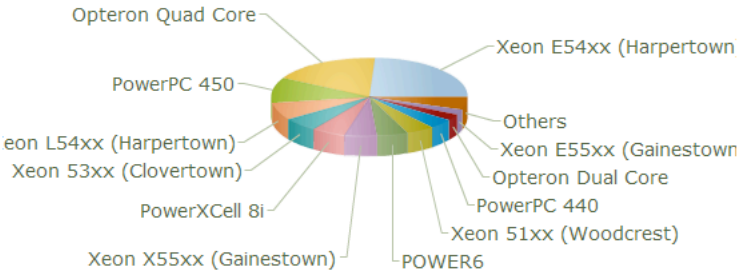

Processor Generation	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
NEC	1	0.20 %	122400	131072	1280
POWER5	2	0.40 %	95630	116128	15280
Xeon EM64T	1	0.20 %	53000	64973	9024
PowerPC 440	12	2.40 %	901453	1123942	401408
PowerPC 970	3	0.60 %	105671	152928	16112
Opteron Dual Core	13	2.60 %	594902	841986	147624
POWER5+	2	0.40 %	41616	49734	6544
Xeon 51xx (Woodcrest)	49	9.80 %	1044529	2031278	181760
Xeon 53xx (Clovertown)	40	8.00 %	1311462	2162052	214288
PowerPC 450	14	2.80 %	2363228	2854900	839680
Itanium 2 Dual Core	1	0.20 %	51441	58982	9216
Opteron Quad Core	28	5.60 %	4048742	5325222	589809
Xeon E54xx (Harpertown)	200	40.00 %	5439988	9391542	831596
Xeon L54xx (Harpertown)	55	11.00 %	1472270	2604678	265572
POWER6	18	3.60 %	1158612	1505203	80064
Xeon X54xx (Harpertown)	8	1.60 %	264019	422350	35196
PowerXCell 8i	4	0.80 %	1250070	1649020	146016
Xeon 52xx (Wolfdale)	3	0.60 %	59788	111439	9632
Xeon 73xx (Tigerton)	3	0.60 %	74108	121743	11664
Itanium2 Montecito Dual Core	4	0.80 %	218057	258150	41200
Itanium2 Fanwood	1	0.20 %	20450	26214	4096
Xeon 32xx (Kentsfield)	1	0.20 %	32800	52032	5420
SPARC64 VII	1	0.20 %	110600	121282	12032
Xeon X55xx (Nehalem-EP)	22	4.40 %	1211317	1497903	131224
GRAPE-DR processor	1	0.20 %	21960	84480	8192
Xeon E55xx (Nehalem-EP)	11	2.20 %	472624	831269	82356
Opteron Six Core	2	0.40 %	67060	81101	8448
Totals	**500**	**100%**	**22607996.30**	**33671604.92**	**4104733**

Distribuição por Fabricantes

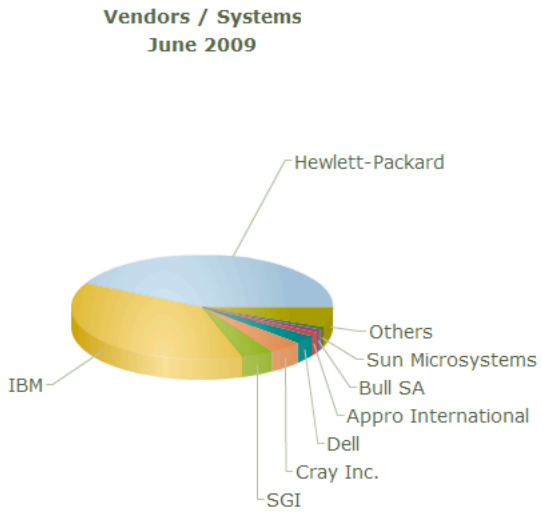

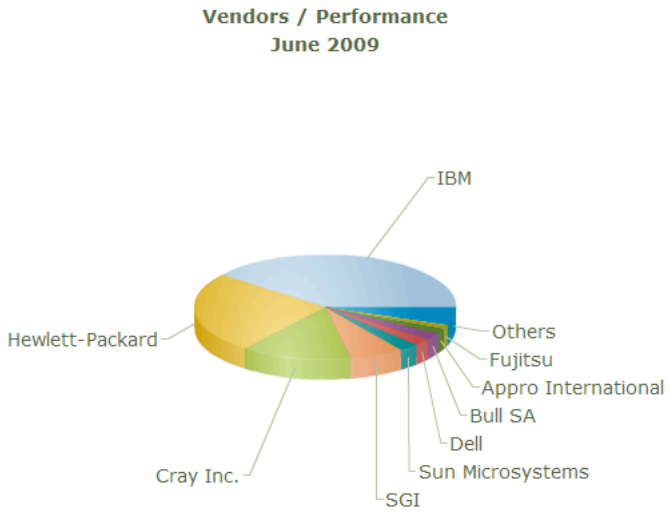

Vendors	Count	Share %	Rmax Sum (GF)	Rpeak Sum (GF)	Processor Sum
Cray Inc.	20	4.00 %	3089355	4002438	455745
Dell	14	2.80 %	492550	674168	72700
Fujitsu	3	0.60 %	249000	279277	26944
Hitachi	2	0.40 %	139634	190055	16384
IBM	188	37.60 %	8903270	13599251	2080672
Intel	1	0.20 %	36000	41714	3408
NEC	2	0.40 %	173190	191283	6656
SGI	20	4.00 %	1504108	1786846	167328
Sun Microsystems	5	1.00 %	549410	729762	80248
Appro International	6	1.20 %	393320	484321	62880
Atipa Technology	2	0.40 %	35080	40128	4180
Linux Networx	1	0.20 %	40610	52992	4416
Megware	1	0.20 %	19390	22270	1856
Self-made	2	0.40 %	41650	119523	11952
Hewlett-Packard	212	42.40 %	5678374	9816965	926948
Dawning	1	0.20 %	180600	233472	30720
Bull SA	5	1.00 %	481080	546316	54576
NEC/Sun	1	0.20 %	87010	163188	31024
DALCO AG Switzerland	1	0.20 %	28990	34560	2880
ClusterVision/Dell	2	0.40 %	37200	52462	4332
Koi Computers	1	0.20 %	37420	56448	6720
Pyramid Computer	1	0.20 %	32800	52032	5420
ACTION	1	0.20 %	38170	49731	5336
UNICORNER/Fujitsu-Siemens	1	0.20 %	18810	24576	2048
ClusterVision/IBM	1	0.20 %	28775	37632	3360
SKIF/T-Platforms	1	0.20 %	47170	60000	5000
LUFAC Computacion SA de CV	1	0.20 %	18480	25440	2120
Raytheon-Aspen Systems/Appro	1	0.20 %	26730	36601	3440
Dell/Sun	2	0.40 %	97020	122189	13224
Lenovo	1	0.20 %	102800	145965	12216
Totals	**500**	**100%**	**22607996.30**	**33671604.92**	**4104733**

Intel Quad Core Nehalem

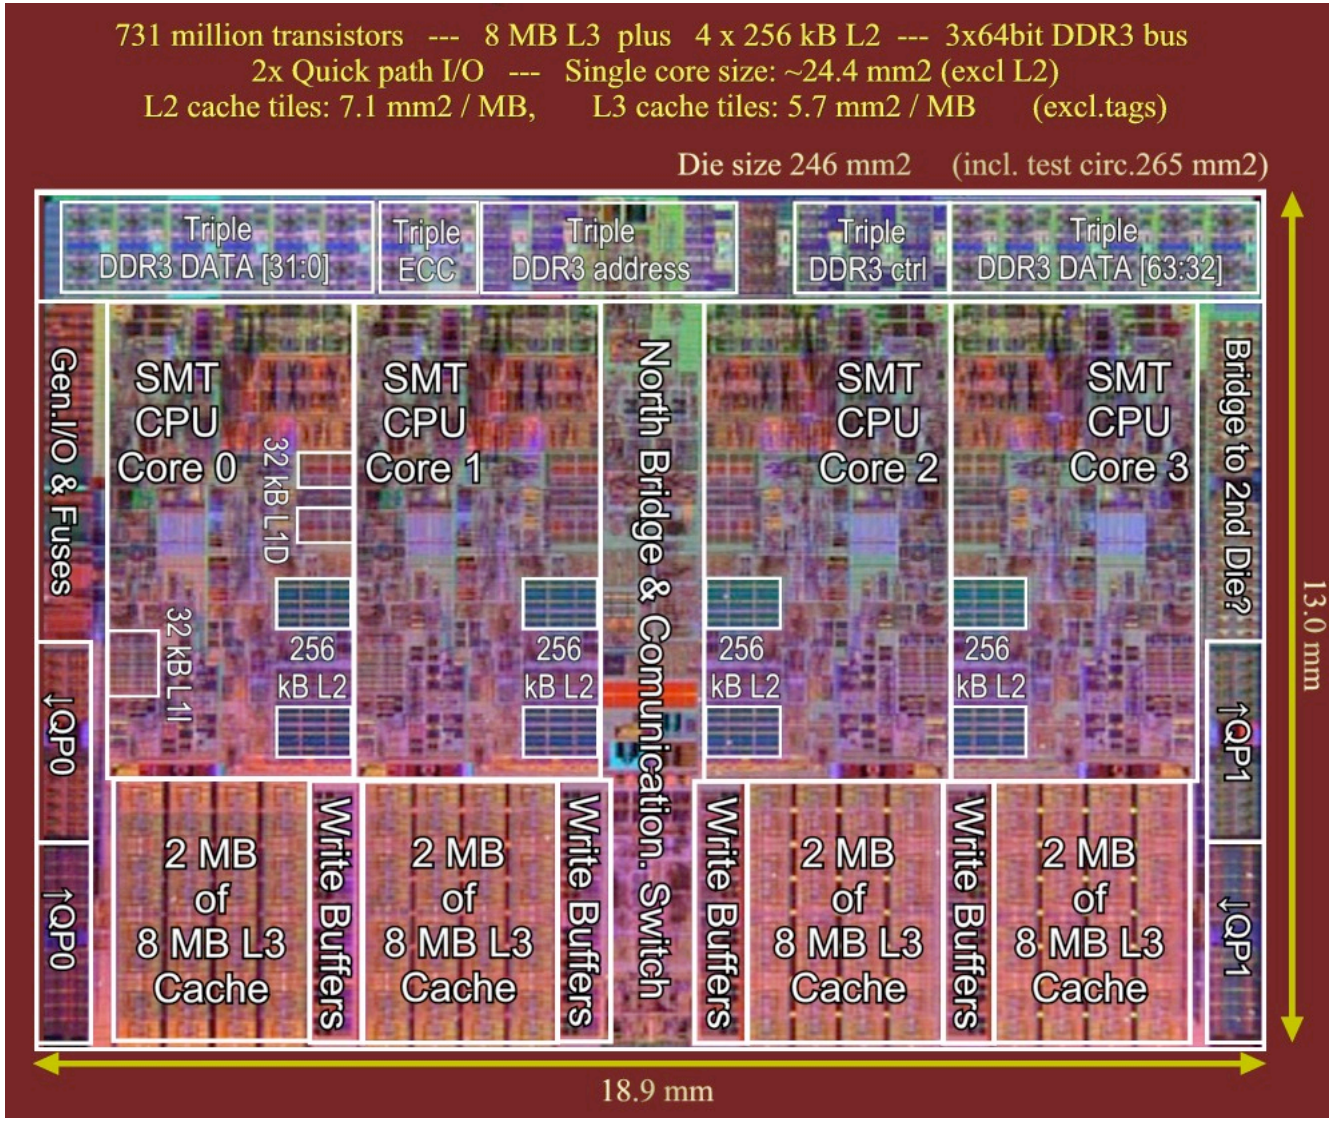

http://www.3dnow.net/phpBB2/viewtopic.php?f=1&t=1474&p=6720

Arquitectura Intel Nehalem

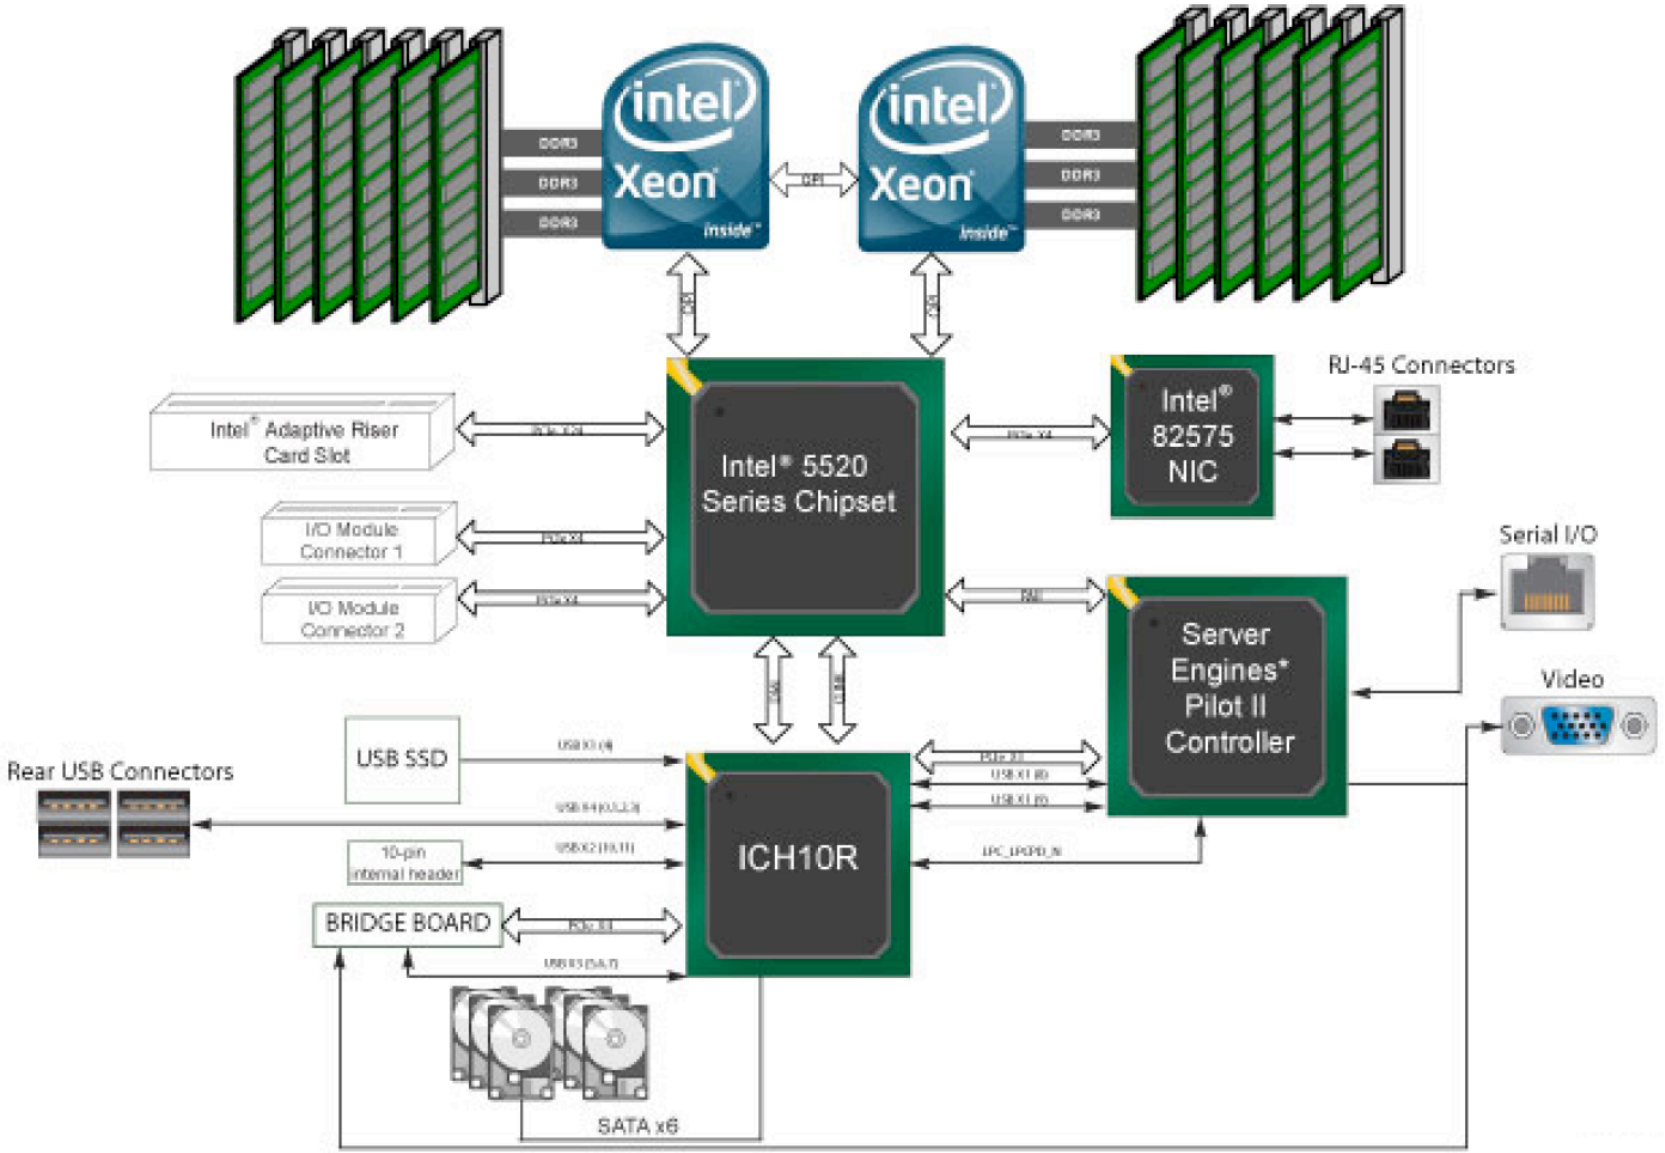

AMD Quad Core Shanghai

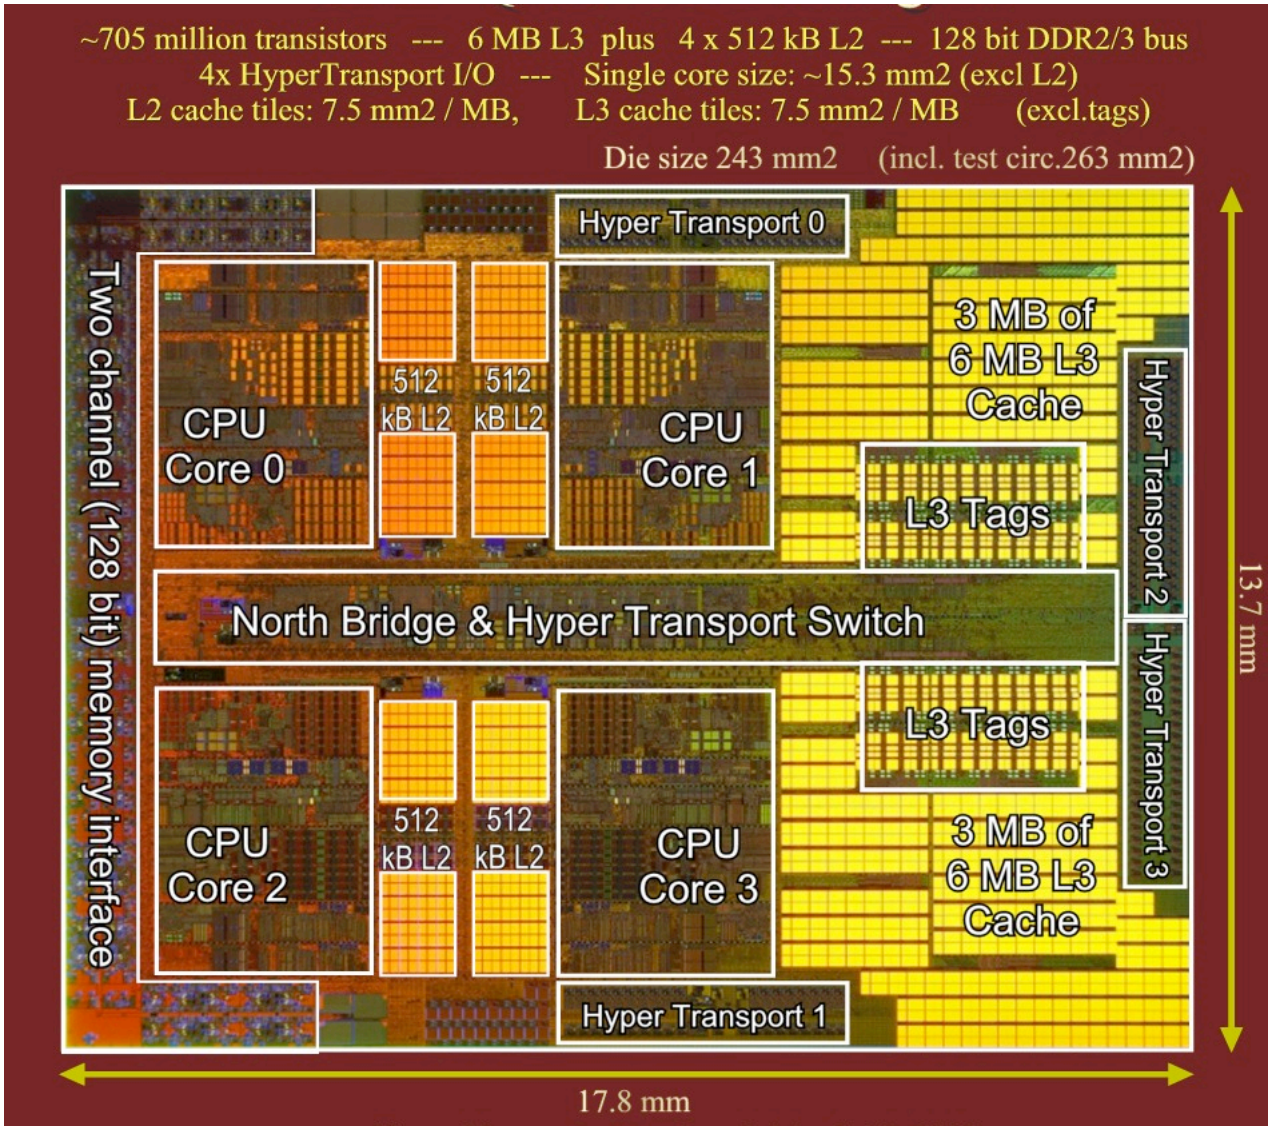

Arquitectura AMD Shanghai

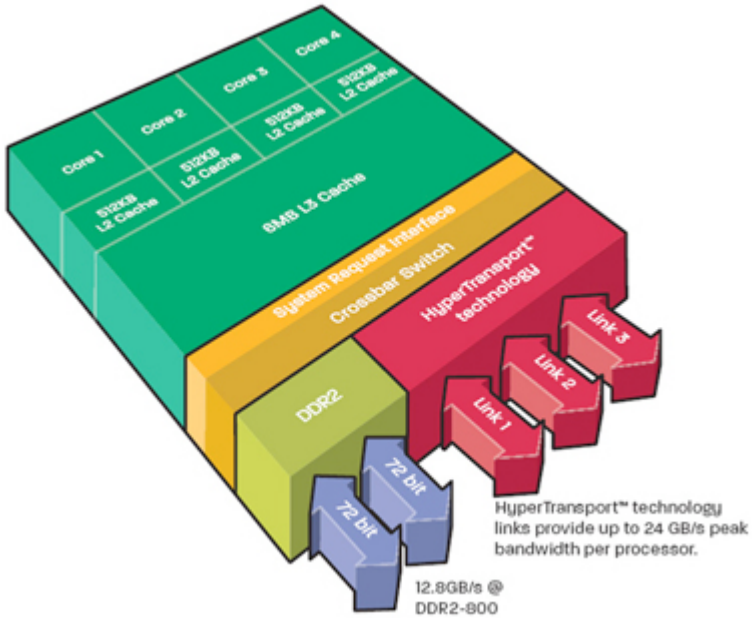

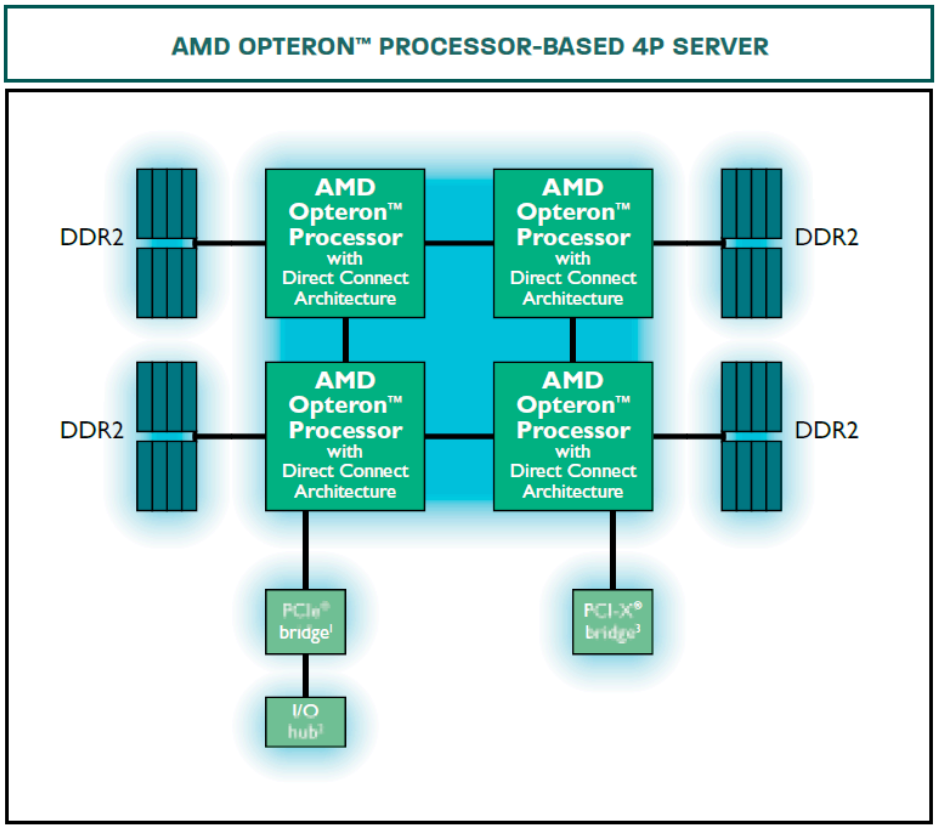

Intel Nehalem-EX Octo-core

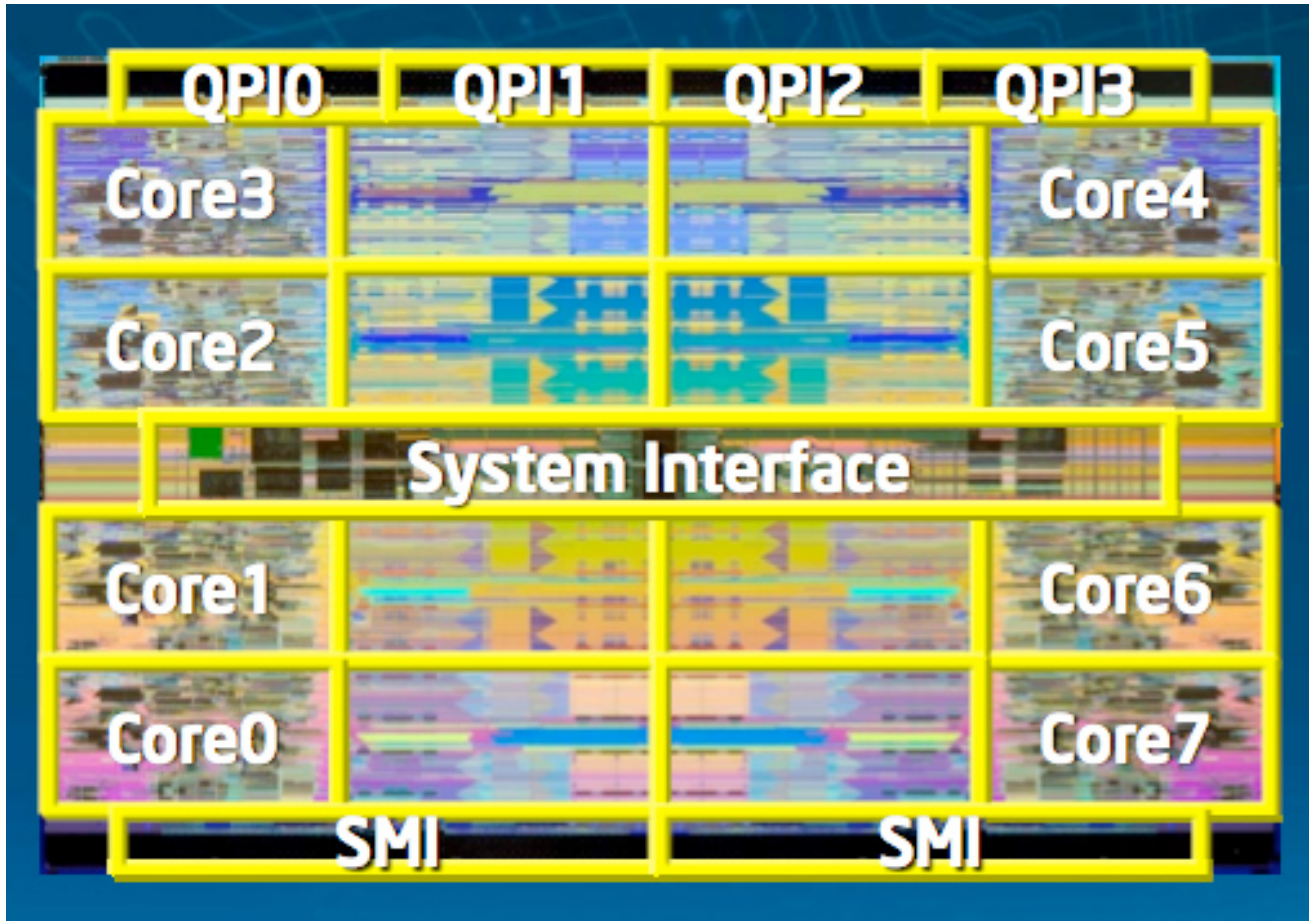

Arquitectura Intel Nehalem-EX Octo-core