

# Linux e la ricerca scientifica

Roberto Ferrari

Parma LUG

Linux Day 2009

24 ottobre 2009

# DISCLAIMER

Titolo molto ambizioso ... non mi allarghero'  
cosi' tanto:

mi limitero' all'ambiente della fisica  
delle particelle (I.N.F.N., CERN)

... ma, per favore, non prendetemi  
troppo sul serio !

(sono solo un fisico ... sostanzialmente uno smanettone)

# Pausa Pubblicitaria

- I'I.N.F.N.

(Istituto Nazionale di Fisica Nucleare)

- il CERN

(European Organization for Nuclear Research)

... cosa ci si fa (al CERN e all'INFN) con i computer

# L'INFN

**ENTE DISTRIBUITO SU TUTTO  
IL TERRITORIO**

19 sezioni, 11 gruppi, 4  
laboratori nazionali,  
1850 dipendenti.

Raggruppa 5000  
ricercatori, la maggior  
parte universitari.



sezione  
gruppo  
laboratorio

**FORTEMENTE INTEGRATO CON  
IL SISTEMA UNIVERSITARIO**

# CERN

Fondato nel 1954 da 12 stati membri (ora sono 20)

Laboratorio di ricerca europeo per la fisica delle particelle → il più grande del mondo nel settore

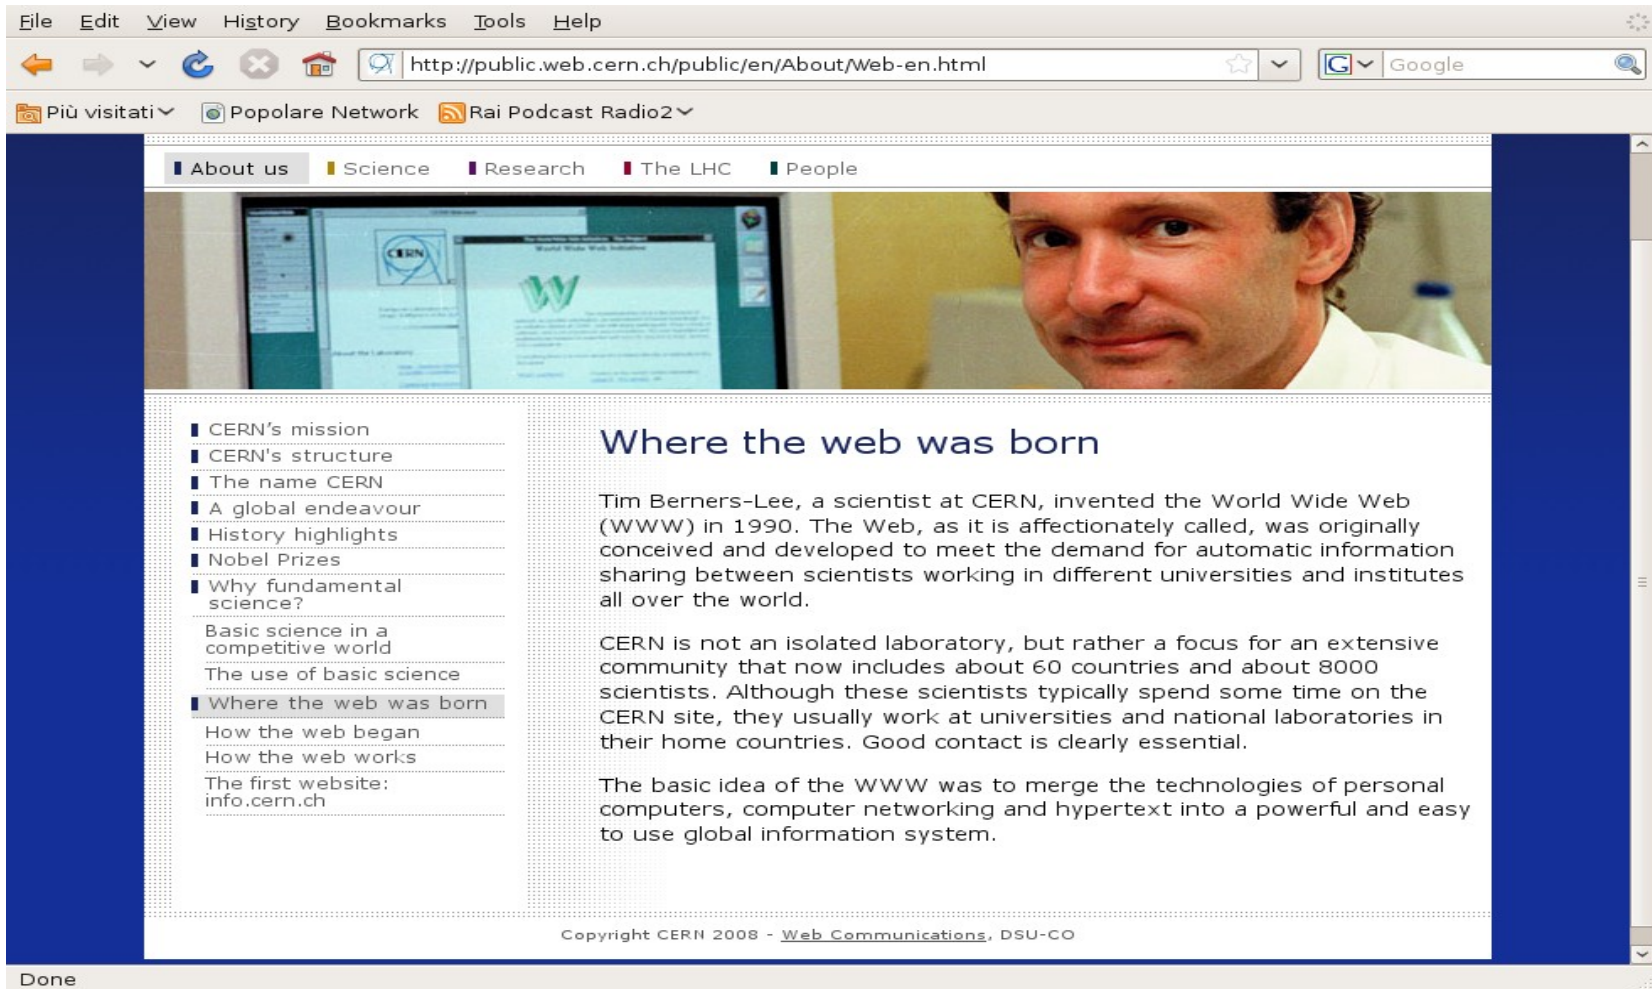
## Dati 2007:

- 2600 dipendenti (staff), dei quali circa 1000 fisici e ingegneri
- coinvolge 9000 ricercatori da 560 istituti in 59 nazioni
- bilancio ~ 600 M Euro (bilancio INFN ~ 270 M Euro)

Stati membri contribuiscono proporzionalmente al proprio prodotto interno lordo (Italia ~ 13%)

Esperimenti: collaborazioni internazionali (Italia -> gruppi INFN)

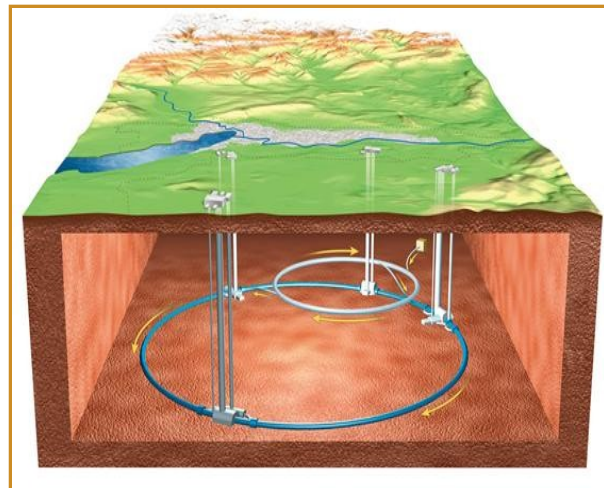
# Dove e' nato il Web ?



ha compiuto 20 anni - <http://info.cern.ch/www20>

# LHC

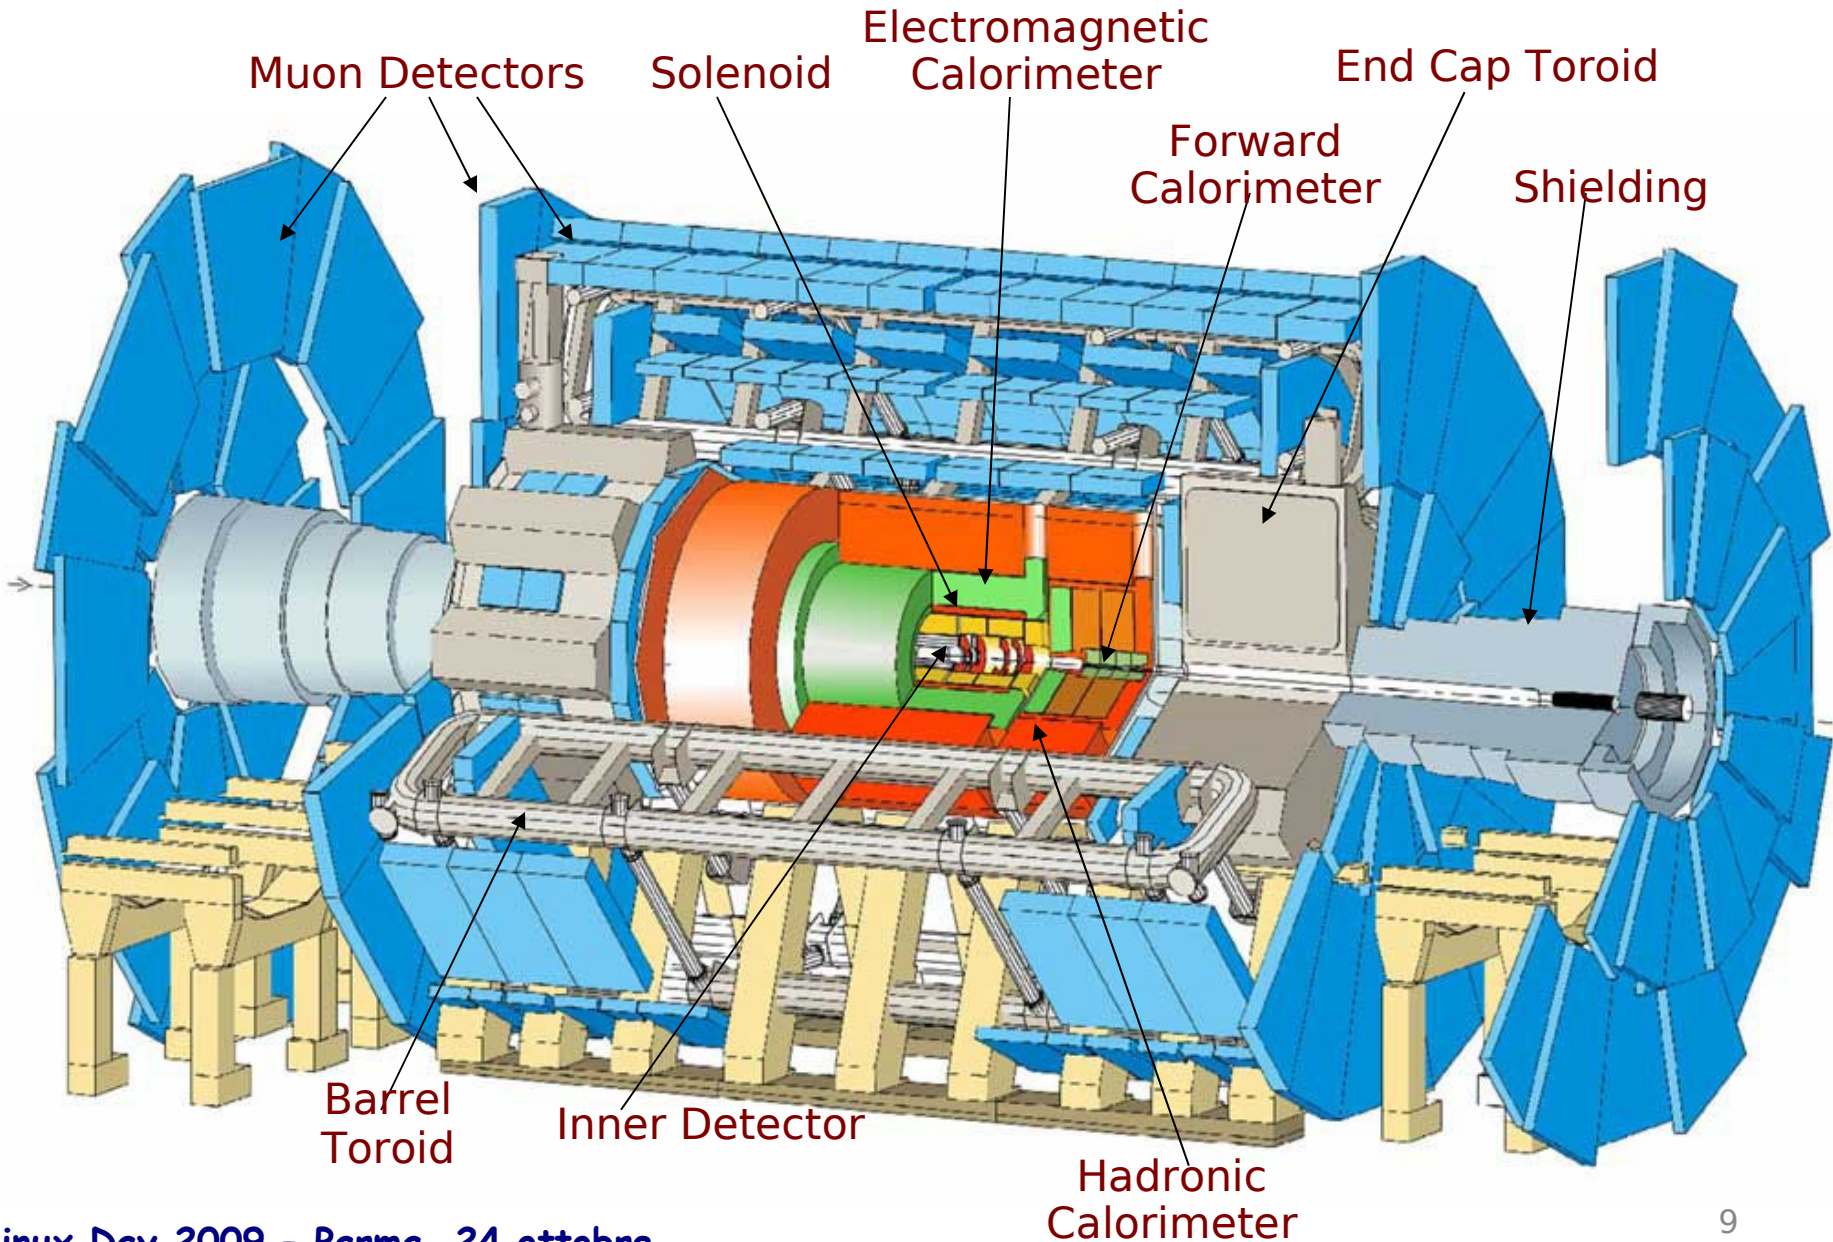
- 27 km di circonferenza
  - (protoni contro protoni) e (piombo contro piombo) ogni 25 ns
  - collisioni a 14 TeV (protone-protone) / 1150 TeV (piombo-piombo)
  - consumo ~ 120 MW (meta' del totale CERN)
- 4 giganteschi apparati sperimentali a ~ 100 m di profondita'







# ATLAS: un microscopio alto 22 m e lungo 46 m



# Il Calcolo

Due settori abbastanza ben distinti:

ONLINE (acquisizione dati - DAQ)

→ efficienza, velocità, robustezza, stabilità

OFFLINE (simulazione, ricostruzione e analisi dati)

→ precisione, ripetibilità

entrambi implementati su più livelli ...

Connessioni (oltre ai dati): descrizione del rivelatore  
(geometria), calibrazioni, condizioni

→ DATABASE

# C'era una volta ...

Anni 80:

minicomputer: DIGITAL con VMS (PDP, VAX, microVAX)

- acquisizione dati, calcolo interattivo, ...

mainframe: IBM 379 con MVS e VM/CMS

- code batch
- grossi "job" di ricostruzione e simulazione dati

supercomputer: CRAY X-MP 48 con UNICOS (!)

- calcolo vettoriale
- UNIX con code batch (sviluppate in casa)
- clock ~ 118 MHz, RAM 128 MB
- potenza di calcolo  $\ll \frac{1}{2}$  Xbox
- costo ~10 M\$

rete: - DECNET, BITNET, X.25  
+ CERN INTERLINK (bridge)

linguaggio: FORTRAN 77

# Qualcosa cambia ...

Fine anni 80 - primi anni 90:

Rete: internet inizia a prendere il sopravvento

Progresso nei compilatori, memorie sempre meno costose:

nascono microprocessori RISC (MIPS, SPARC, ALPHA, ...)

molto aggressivi e competitivi ... UNIX-based

Si abbandonano i grossi mainframe a favore di un approccio distribuito a bassi costi unitari (mini/micro computer).

Parole chiave: RISC, UNIX, TCP/IP, SCSI, C, \* **SCALABILITA'** \*

Prima macchina: Apollo DN10000 (CERN 1989)

Soluzioni eterogenee SGI, SUN, Apollo, HP, DEC, ... IBM

... ognuna con un suo UNIX proprietario

# Babilonia di Dialetti

SGI: microprocessore MIPS con sistema operativo IRIX (1)

SUN: SPARC con SunOS4 (2) e poi SOLARIS (1)

Apollo: Domain/OS (3)

HP: PA-RISC con HP-UX (1)

DEC: ALPHA con Ultrix (2) e poi Digital Unix (3)

IBM: RS/6000 e PPC con AIX (1)

(1) ~ System V

(2) ~ BSD

(3) ~ ibrido

(comunque molto piu' semplice di prima ...)

... fine anni 90 - primi anni 2000 inizia la transizione a Linux:

prima Red Hat Linux (!) poi Scientific Linux Cern

# S.L.C.

Scientific Linux: release creata e mantenuta da FermiLab e Cern (piu' altre universita' e laboratori nel mondo)

Nasce nel 2004 a Fermilab

"Red Hat Enterprise Linux" ricompilata e integrata con pacchetti specifici:

<https://www.scientificlinux.org/>

Scientific Linux Cern: sottovariante CERN

<http://linux.web.cern.ch/linux/scientific.shtml>

Nel marzo 2009, circa 36000 sistemi girano Scientific Linux (SL) e circa 14000 girano Scientific Linux CERN (SLC)

+ L.S.F. (gestione code batch), CASTOR (gestione storage) ...

# Acquisizione Dati (DAQ)

Dalla ~ meta' degli anni 80 (LEP):

Da VAX/VMS → Single Board Computer (SBC) VME

- processori 680x0
- sistema operativo real time OS-9

\* REAL TIME O.S. : esiste un ritardo massimo di risposta ben definito \*

Il kernel UNIX "standard" non e' real time:

una chiamata di sistema puo' richiedere un tempo indefinito

# DAQ .vs. LHC (ATLAS) ...

Inizi anni 90 - R&D per LHC:

parole chiave: VME + RISC + UNIX real time + \* SCALABILITA' \*

es: board (SBC) MIPS R3000/R4000, PowerPC con LynxOs

Seconda meta' anni 90:

front-end (SBC): LynxOs + ppc (ma anche test con windows NT)

back-end: Solaris/HP-UX/... (poi primi pc con linux)

la richiesta di "real time" cala rapidamente ...

rimangono delle diffidenze verso il mondo "open source"

Ultimi 10 anni, graduale convergenza verso:

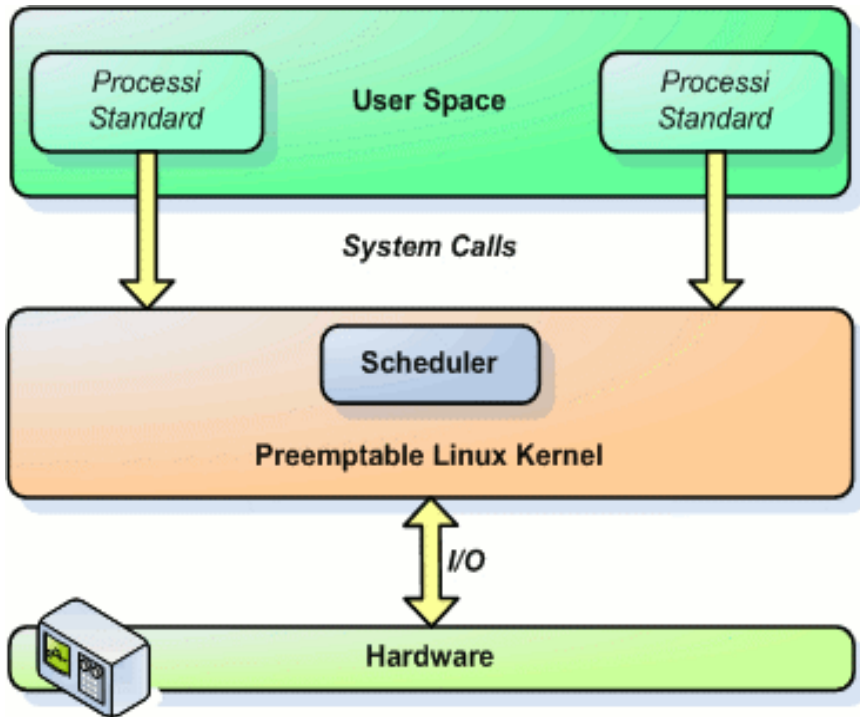
Red Hat Linux !

front-end (SBC): 80x86 + linux (Red Hat Linux → S.L.C.)

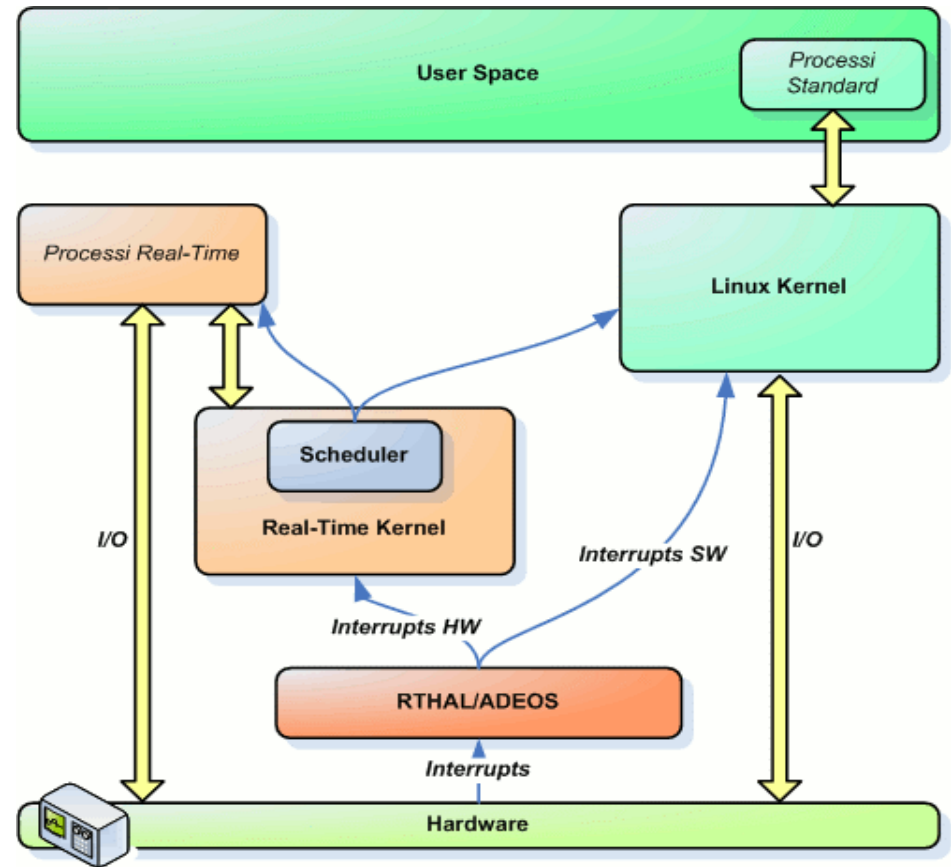
back-end: rack di macchine linux (Red Hat Linux → S.L.C.)



# UNIX Real-Time



Low-latency patch (Ubuntu Studio):  
linux kernel interrompibile



RTAI: il kernel linux gira come una  
applicazione con priorit  maggiore

# DAQ ATLAS

~ 40 M di eventi / sec

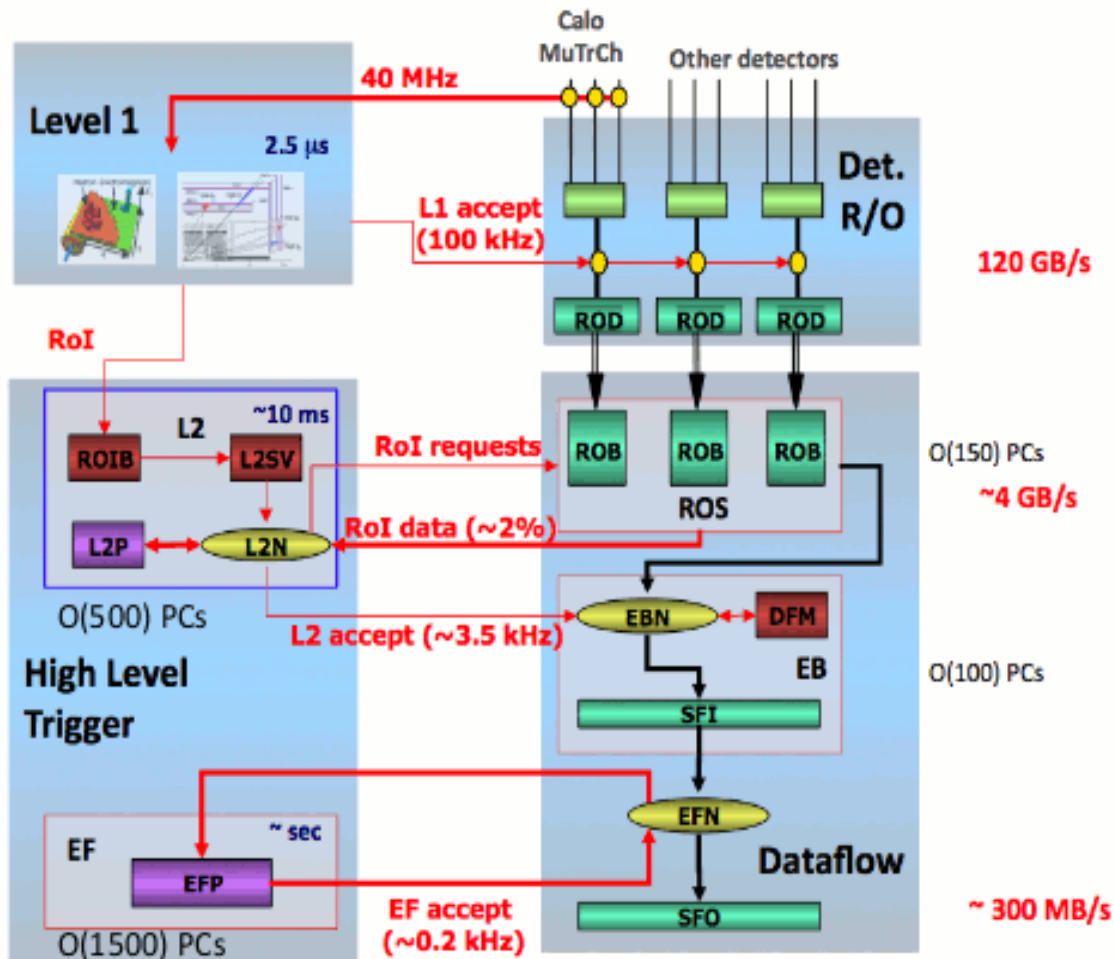
~ 1 evento ogni  $10^9$

\*\*\*INTERESSANTE\*\*\*

~ 100 M di segnali → 1.5 MB/evento

## Selezione eventi "on-line"

- Elettronica e computer dedicati
- migliaia di processori in parallelo (hardware)
- decine di migliaia di processi da controllare (software)



# Lista della spesa ...

~ 100 rack x ~ 30 macchine = ~ 3000 macchine "rack mounted"

in generale dual cpu / quad core / 16-24 GB

delle quali ~ 2300 per processamento dati (calcolo)

~ 100 file server (normalmente 1 per rack)

~ 150 trasferimento dati

~ 300 per monitoraggio, controllo, servizi vari

inoltre ~ 160 SBC (VME) + 150 ROS (readout system)

## Performance:

1 macchina (8 core) : 200 eventi/s al livello 2 (40 msec/core),

2 eventi/s all'event filter (4 sec/core)

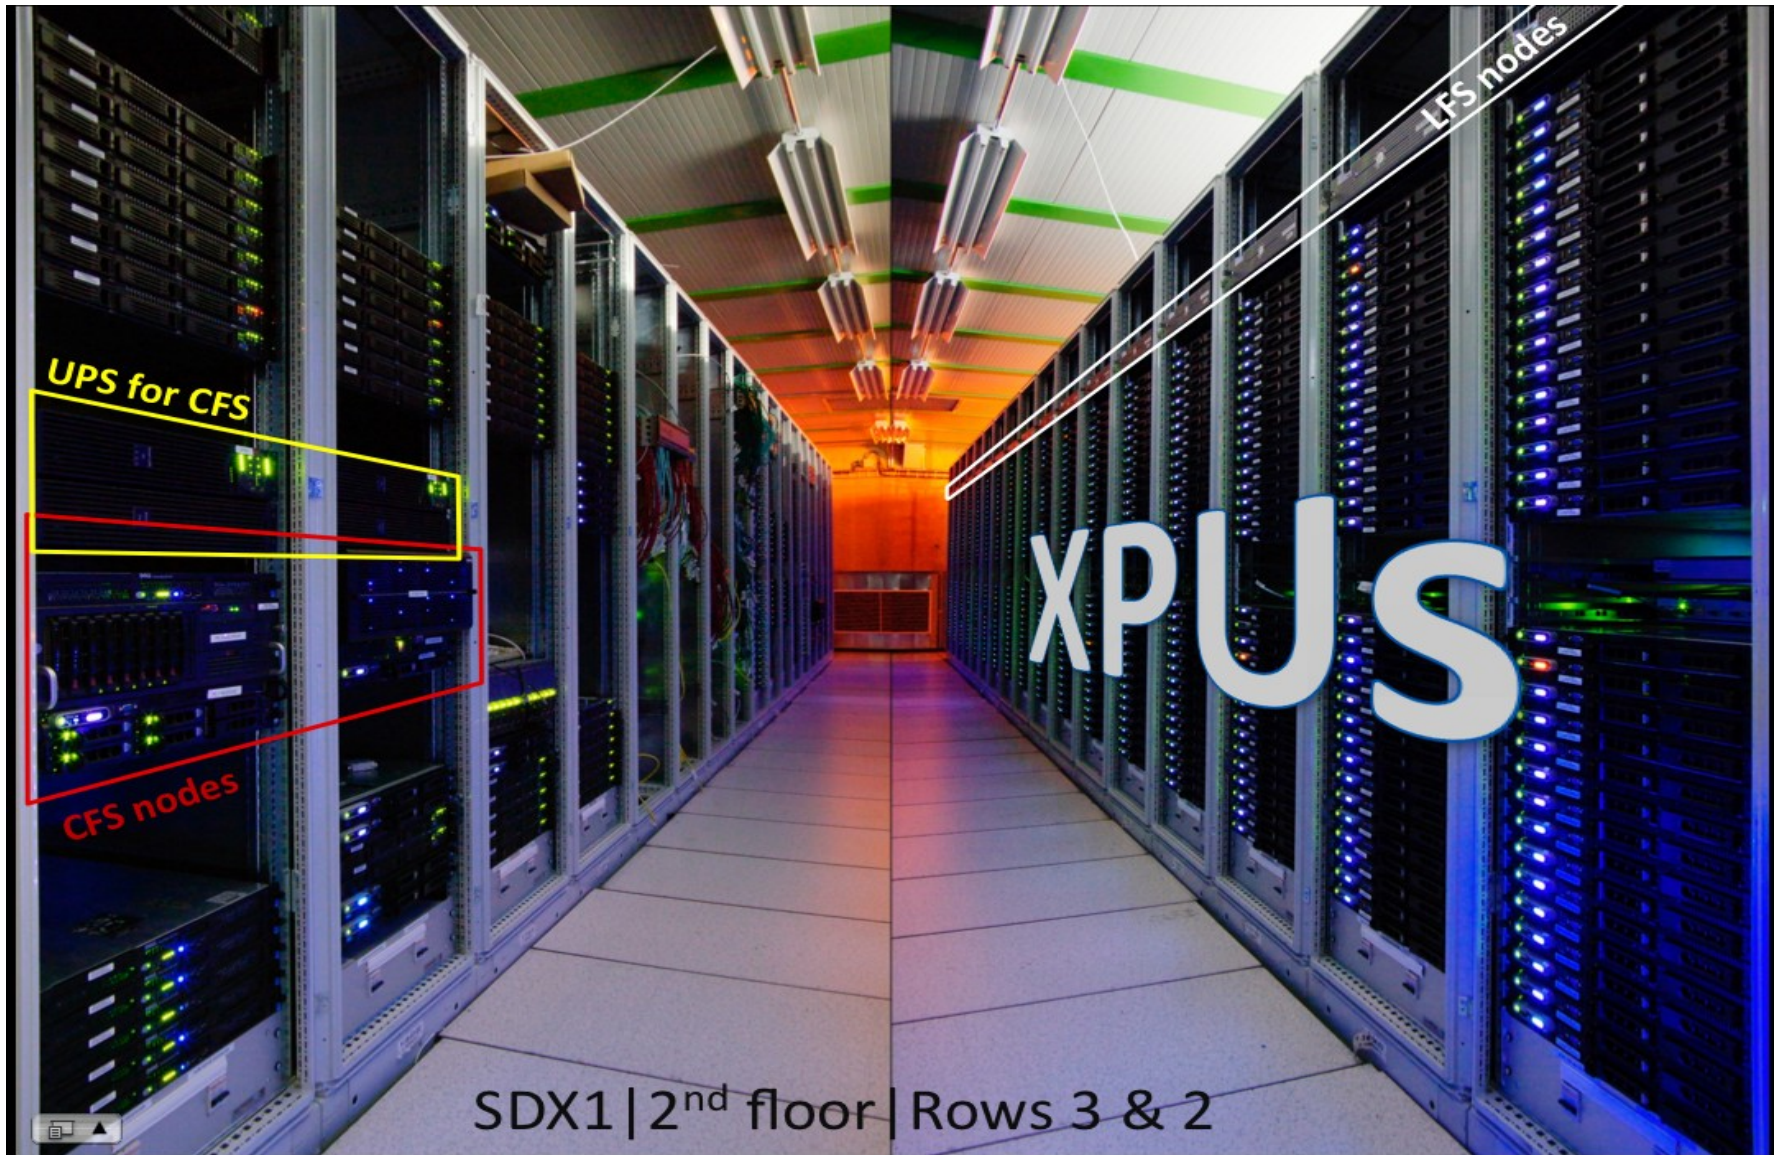
Storage (cache): 6 macchine 6U x (24 x 0.5 TB) = 6 x 12 TB (raid5) = 72 TB

E4 Computer Engineering (Scandiano):

24 TB (24 disch x 1 TB), 2x4 core (16 "processori" ind.), 24 GB

Link verso il centro di calcolo: 2 x 10 Gb/s

# I Rack



# La Sala di Controllo



# Software ...

Trasferimento, processamento, monitoraggio dati ("DataFlow"):

C/C++ (protocolli di rete: UDP, TCP)

GUI: Java / Qt / Python (tk/tcl)

Sistema Esperto: Common Lisp

Inter Process Communication: CORBA

Configurazioni/Calibrazioni/Allineamenti/Geometrie:

file, OKS (xml), COOL, ORACLE, SQLITE, Python ... largo uso di Proxy

Documentazione: WWW, Twiki

... Nagios (monitoraggio !), IPMI (controllo !) ...

Parole chiave: Macchine a Stati Finiti, Scalabilita',

Partizionabilita', Configurabilita', Sicurezza

DAQPanel (on pc-atlas-cr12.cern.ch)

Insert Here Some Info

Setup Script: /sw/tdaq/setup/tdaq-02-00-03.sh

Part Name: ATLAS

Database File: /atlas/oks/tdaq-02-00-03/combined/partitions/ATLAS.data.xml

Buttons: Start Partition, Monitor Partition, RC Status, Local Procs, OKS, DVS, Log Manager, MRS, Busy, DQM Display, Trigger Presenter, Event Dump, OHP, OMD, ISPY, SFO Display, Get Default, Read Info, Get Partition.

Log Messages

Log Messages window content:

```

ohp options --> <
ohp options --> <
TriP options --> <
TriP options --> <
BUSY options --> <
OMD options --> <
MRS filter --> <
  
```

Buttons: Resize, Clear Log, Change file, Exit

You are robertof and your role is

# Run Control Macchina a Stati Finiti

ATLAS TDAQ Software Graphical User Interface - Expert Control

RELOAD CONFIGURATION

Run control: RUN CONTROL STATE: RUNNING

START/STOP FLOW

PROCESSES RUNNING. Should coincide with the RUN CONTROL STATE

ERROR LOGGER. Messages for experts so far...

Table of processes:

Process Name	Status
ArchivingServiceImpl	RUNNING
ScfApiDDCServer	RUNNING
ScfApiServer	RUNNING
ScfConfigurationServ	RUNNING
CalibrationController	RUNNING
ScfEndcapASegmentC	RUNNING
ddc_sct_c11	RUNNING

Table of error logs:

Time	Severity	Component	Message
16:12:55	WARNING	ArchivingService...	ScfctisException
16:12:44	WARNING	ArchivingService...	ScfctisException

MRS Monitor [ATLAS]

Partition: ATLAS

Subscription criteria: WARNING (selected), ERROR, FATAL, INFORMATION

TIME	SEVERITY	APPLICATION	NAME	MESSAGE
15:39:41	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3599... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:41
15:39:40	WARNING	LVL2-L2-2-rack...	gatherercissue	Histogram'L2PU-5920... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3752... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3176... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3560... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	ERROR	CheckBCIDGnam	bcidcheck:AnyError	Run 136207 Ev 18434 Ref 1 L1 Ox0c0357 TT Ox0C BC Ox576 Status Ox 1 full 0 - event format error -- 95 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:40	ERROR	CheckBCIDGnam	bcidcheck:AnyError	Run 136207 Ev 18435 Ref 1 L1 Ox0c0357 TT Ox0C BC Ox5ae Status Ox 1 full 0 - FORWARD_BCM, module=12 (opt=0) (ROB Ox81000c) BCID internal mismatch: 0xbdb1 / 0xae -- 95 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:39	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3690... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:40	WARNING	ROS-TILC-LBC-01	ROS:CoreException	Timeout: in request for fragment with L1 ID 2191260819 -- 2252 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3194... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3199... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3162... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:39	WARNING	LVL2-L2-1-rack...	gatherercissue	Histogram'L2PU-8103... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-2-rack...	gatherercissue	Histogram'L2PU-6107... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3208... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	SFI-32	SFEDataIntegrity	Problem with data integrity. Event fragment from ROB Ox81000c (ROS_BCM_ROS_SubDet: 129) with LVL1ID: Oxae040e and BCID: 1357 has a BCID mismatch: Event_BCID - ROD_BCID = -195. [ ROS Fragment status= 0x1 ] -- 57 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:38	WARNING	LVL2-L2-2-rack...	gatherercissue	Histogram'L2PU-6307... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:38	WARNING	LVL2-L2-4-rack...	gatherercissue	Histogram'L2PU-3119... can not be summed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38

Buttons: Clear, Message format, Number of visible rows: 2,000, Current MRS subscription: WARNING/ERROR/FATAL

ATLAS ATLAS - Konqueror (on pc-atlas-cr02.cern.ch)

Location: https://pc-atlas-www.cern.ch/elog/ATLAS/ATLAS/

Electronic logbook for the ATLAS experiment. Page 1 of 2669

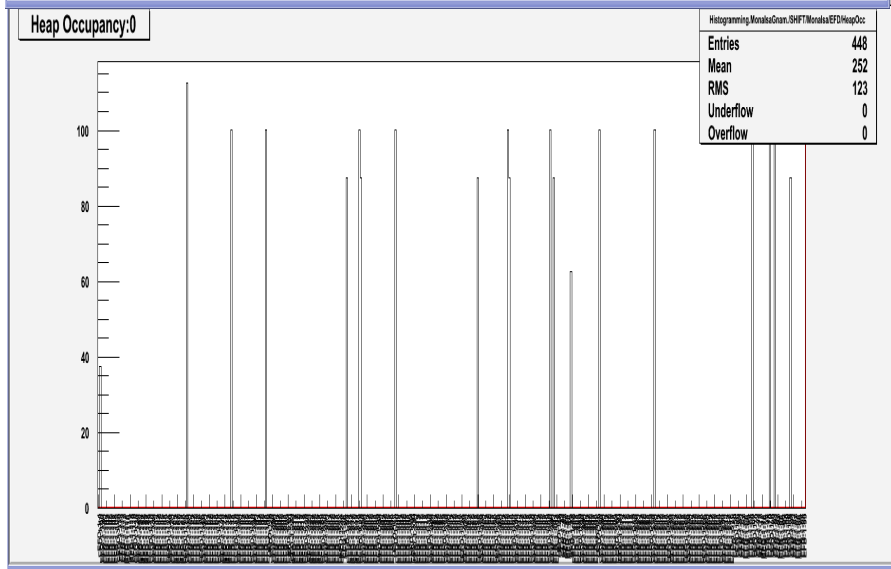
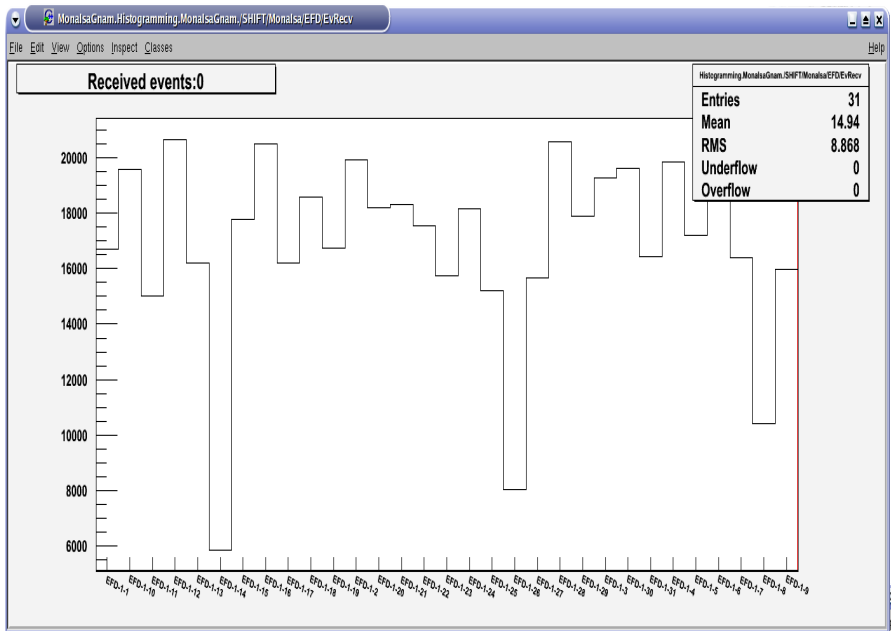
Logged in as "Ferrari Roberto"

5375 Entries

Date	Author	Message Type	System Affected	Subject	Text
23.10.09 15:29	Canepa Anadi	Data Quality		Online DQ Shifter Summary	Summary of online dq shifter Atlantis and VPI running fine
23.10.09 15:26	DCS JS	Slimos-TI	DSS   Tech. Infra.	MAG_Toroid_SlowDump	DOB Alarms: MAG_Toroid_SlowDump
23.10.09 15:17	Tikhomirov Vladimir O	Default Message Type	TRT	New shift	Combined cosmic run #13620 is going thrm. TRT status OK.
23.10.09 15:14	Bondoli Bond	Shift Summary	DAQ	Shift Summary for Run Control desk	run 136183 from previous shift
23.10.09 15:11	Hayakawa Takashi	Default Message Type	TGC	TGC on-call shift report	*** Readout and Trigger timing *** run#136176 run#136183
23.10.09 15:10	Perez Reale Valeria	Shift Summary	TRT   TDI   Cryo   DCS   Pixel   Magnets   TGC   ID Gen. (IC)   MDI   DSS   SCT   LArG Network   RPC   CSC   DAQ   HLT   LVL1   Monitoring   GAS   SysAdmins   BCM   LUCID   Counting Room   Tier0   Beam Conditions   ZDC   Event Displays	Shift Summary for Shift Leader	07:00 run from started at 4am last night is #1311 running: run 136183 (620, 881, 775) / (R) . Initially.
23.10.09 15:06	Kim Tae Jong	Shift Summary	HLT   LVL1	Shift Summary for Trigger desk	Run upon arrival: #136183 Going since 04:00
23.10.09 15:05	Qi Ming	Shift Summary	TGC	Shift Summary for Moon Desk 3 - TGC	Shift summary: 07:00-15:00, 23 Oct. 2009
23.10.09 15:04	Dubbert Jorg	Default Message Type	DCS   MDT	BOG6A12 MLI HV interlock asserted	Asserted HV interlock for BOG6A12 MLI chamber triggered repeatedly since yesterday. Stable for some time (up to some hours) then East
23.10.09 15:03	Ghodbane Nabli	Title	Title	end of shift summary	Opening run: 136207 List of good runs with L1cal in : Summary of what happened during the shift.
23.10.09 15:02	Ferretti Claudio	Shift Summary	TGC   MDT   RPC   CSC	Shift Summary for Moon Desk 1 - MDT/CSC	Many runs (staged mostly because of AC problems) 136183 - 160 Rev. Run 136183 ended at about 9:00. Bectors In

# Monitoraggio Online

# Information Service



Partition 'ATLAS', server 'DF-EF-Segment-01-rack-Y03-0602-iss'

Name	Type	Modified	Description
EFD-1-25	EFD	16/7/08 09:43:31,549965	
EFD-1-26	EFD	16/7/08 09:43:34,503773	
EFD-1-27	EFD	16/7/08 09:43:31,834124	
EFD-1-28	EFD	16/7/08 09:43:31,946579	

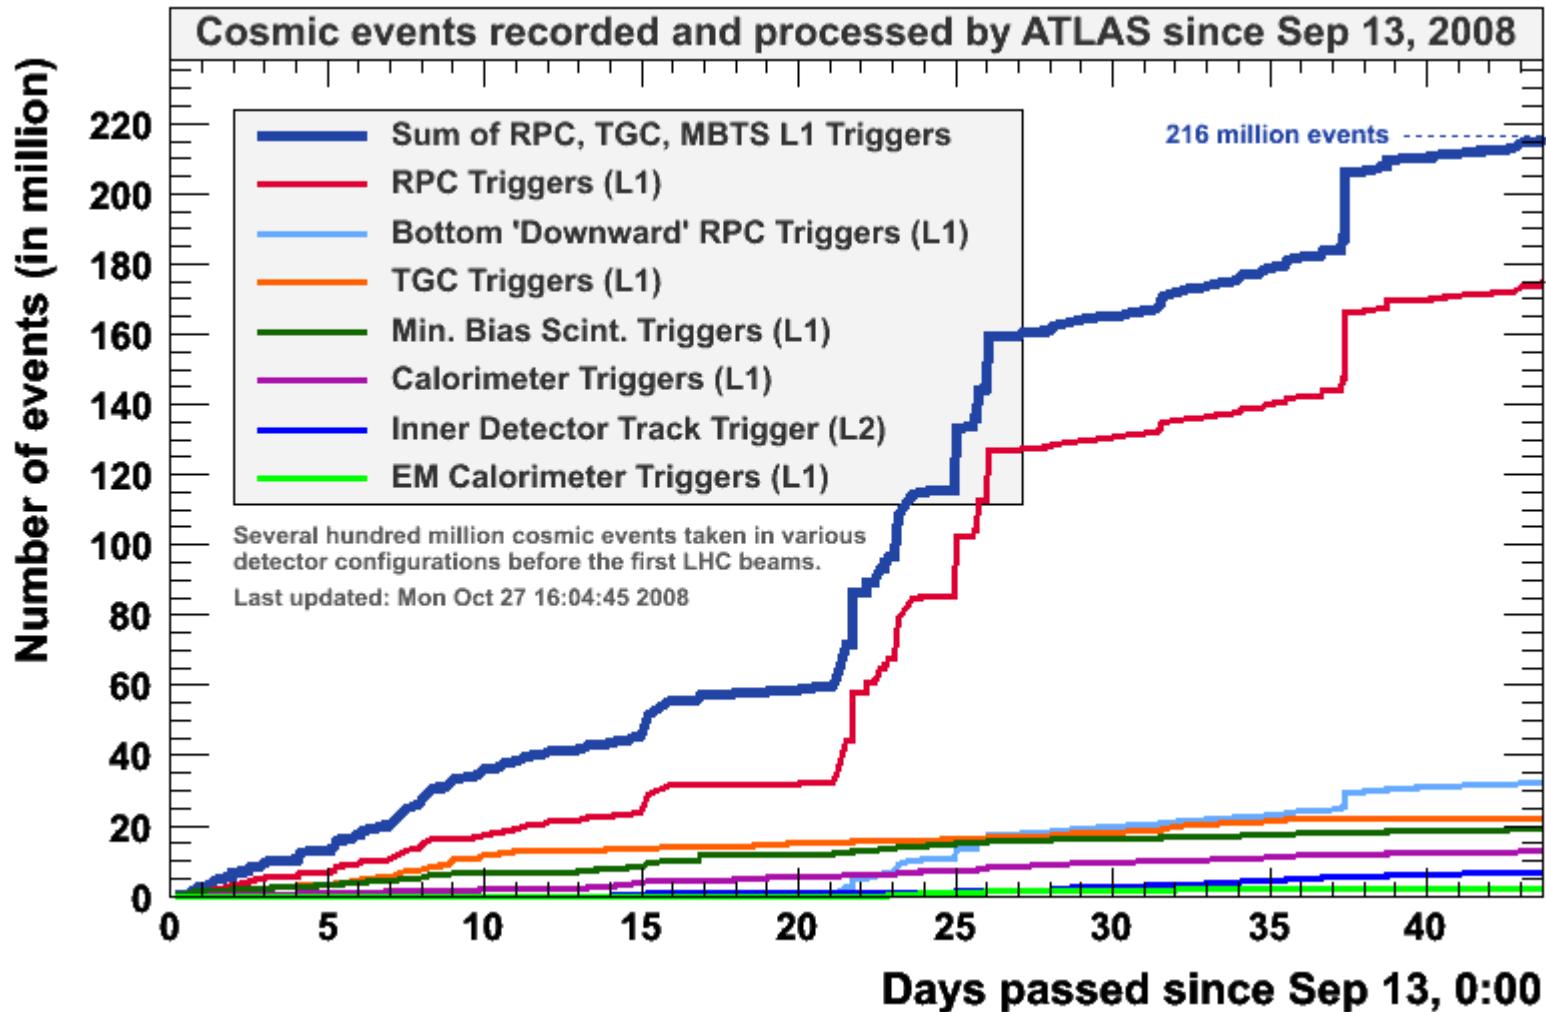
  

Value	Type	Name	Description
pc-tdq-xpu-0245:/local_L/efHeap/sharedHeap.cmc.ATLAS	String	SharedHeap	SharedHeap file fullpath
3	UI16	ConnNbrSFIs	Number of connected SFIs (sum over InputTasks)
5	UI16	ConnNbrSFOs	Number of connected SFOs (sum over OutputTasks)
4	UI16	ConnNbrPTs	Number of connected PTs (sum over ExtPTsTasks)
87.54	Double	HeapOcc	SharedHeap occupancy (%)
1521	S32	EventsRcv	Number of received events
1514	S32	EventsSent	Number of events sent to SFO (ie: Dismissed-Deleted)
7	S32	EventsInside	Number of events inside
0	US2	EventsWaitingForProc	Number of events waiting for processing
2	US2	EventsWaitingForDeli	Number of events waiting to be sent to SFO
0	Double	RateIn	Current rate of incoming events (Hz)
0	Double	RateOut	Current rate of events sent to SFO (Hz)
0	Double	FluxIn	Current rate of space allocation in SH; >M data flux
0	Double	FluxOut	Current data flux to SFO (MB/s)
-1	Double	FlowCtrlStopTime	Guess of the stop transition time (s)
460	US2	FlowCtrlISleepTime	Current flow control sleep time (ms)
538	US2	FlowCtrlBarrierLocks	Number of times the input barrier has been locked
0	S32	ptionNbrProcTimeouts	Number of processing timeouts
0	S32	ptionNbrSocketHungUps	Number of PT socket hung-ups
0	S32	ptionNbrForceAccept	Number of force accepted events
0	S32	efionNbrSFiBrokenConn	Number of broken connections to SFI
0	S32	efionNbrSfoBrokenConn	Number of broken connections to SFO
1521, 0, 0, 0, 0, 0	S32[6]	EventTagTypesIn	Type counters: phys, calib, reserved, debug, unkno
1519, 0, 0, 2, 0, 0	S32[6]	EventTagTypesOut	Type counters: phys, calib, reserved, debug, unkno

403 objects | 24 attributes



# Dati 2008 (1 evt = 1.5 MB)



# Offline

$O(1 \text{ miliardo})$  di eventi all'anno da ricostruire e analizzare  
~ Altrettanti da simulare

## STORAGE

$1.5 \text{ MB/ev} \rightarrow 1.5 \text{ PB/anno} \times \sim 2 = \sim 3 \text{ PB/anno}$  di storage

## CPU

processamento:  $\sim O(50) \text{ kSi2k*s/evt}$  (2-3 volte l'anno)

simulazione:  $\sim O(100) \text{ kSi2k*s/evt}$

Totale (solo ATLAS):  $\rightarrow \sim 220 \text{ G kSi2k*s} \sim 7000 \text{ kSi2k*anno}$  (eff.)

a cui vanno aggiunte risorse per calibrazioni, test, analisi finale

# (Parentesi)

## Analisi Eventi

Dalla DAQ arrivano informazioni "grezze" (misure di tempi, cariche elettriche, tensioni)

Ricostruzione a piu' stadi (attivita' centralizzata):

→ informazioni fisiche (posizioni, velocita')

→ identificazione particelle, energia, quantita' di moto

Analisi fisica (attivita' caotica):

→ criteri di separazione fondo / segnale (selezione eventi)

→ analisi statistica

Lo stesso esercizio viene fatto con le simulazioni (la sola differenza rispetto ai dati reali, e' che ogni evento contiene anche la "verita' montecarlo")

# Simulazione, Ricostruzione e Analisi Dati

Attività distribuita su più livelli:

Tier-0 (CERN) → Tier-1 (grossi centri nazionali)

→ Tier-2 (centri regionali) → Tier-3 (istituti)

Dati distribuiti con ridondanza (almeno due copie di ogni dataset)

Cataloghi (database) per tenere traccia

Esecuzione delocalizzata: nuovo strato software (middleware) che indirizza gli eseguibili dove si trovano i dati, raccoglie e assembla i risultati

## LA GRID !

(N.B.: il paradigma della Grid è ancora più forte)

# La Griglia (GRID)

Dati LHC equivalenti a ~20 milioni di CD (una pila alta 20 km) all'anno

Per l'analisi necessari ~100mila dei più veloci processori odierni



WWW: accesso a informazione archiviata in diverse località geografiche

GRID: accesso a risorse di calcolo e di archiviazione dati distribuite su tutto il pianeta



# In Italia

Tier-1: CNAF (Bologna) unico per tutti gli esperimenti LHC (e non solo)

Tier-2: ~10 (Roma, Legnaro, Torino, Napoli, Catania, CNAF, Pisa, Milano)

Investimento (ad oggi) ~ 30 M Euro (incluse infrastrutture CNAF)

+ molti anni uomo di sviluppo sw (anche grazie a finanziamenti europei)

# Il Middleware

## Organizzazioni Virtuali (nel caso ATLAS)

user → GANGA (frontend) → GLITE-WMS (backend)  
→ Risorse fisiche (GRID)

\*\* In funzione delle risorse fisiche, l'implementazione degli stadi intermedi puo' essere diversa (le GRID sono piu' di una) \*\*

“Ganga allows for the specification, submission, bookkeeping and post-processing of computational tasks on a wide set of distributed resource”

## Workload Management System (WMS):

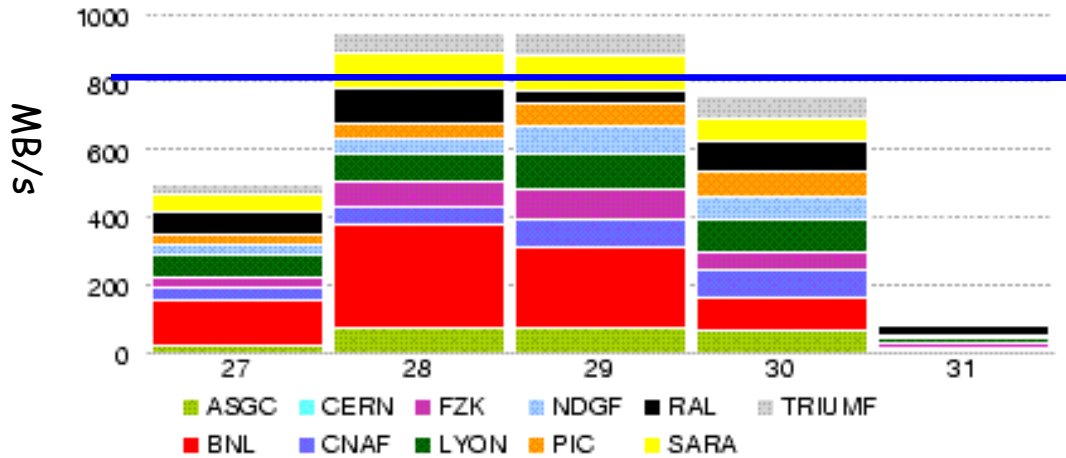
un insieme di componenti del Grid middleware responsabili della distribuzione e gestione di processi attraverso piu' risorse Grid

Al livello hardware:

Allocazione ~ dinamica delle risorse di calcolo (CPU)

Allocazione ~ statica dello storage

# Esiste !

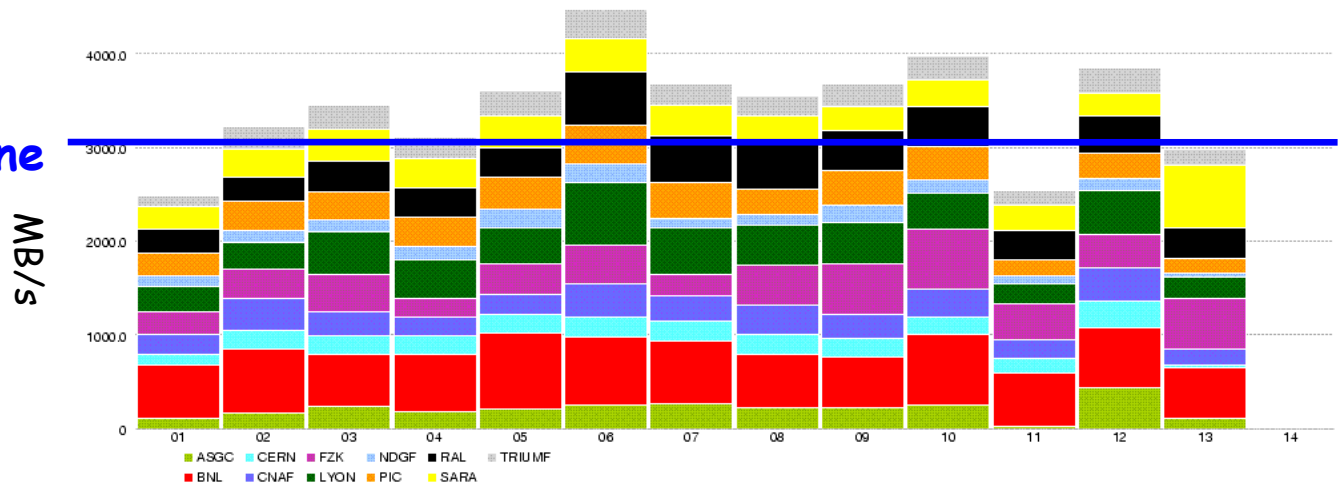


0,8 GB/s 2 giorni

**CCRC08**

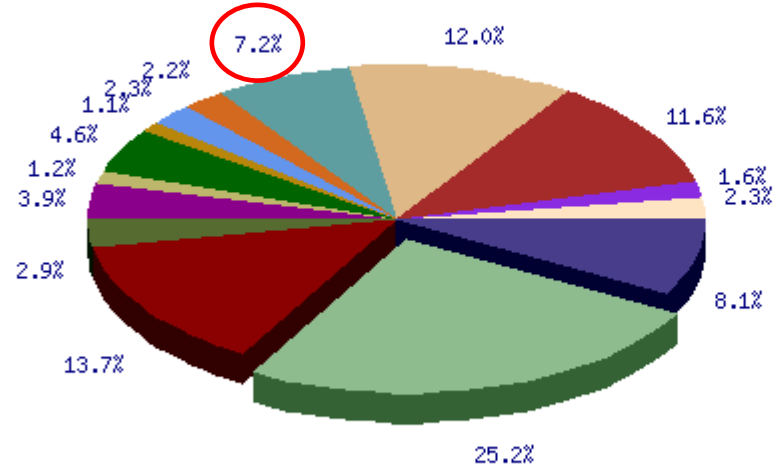
3 GB/s 2 settimane

**STEP09**



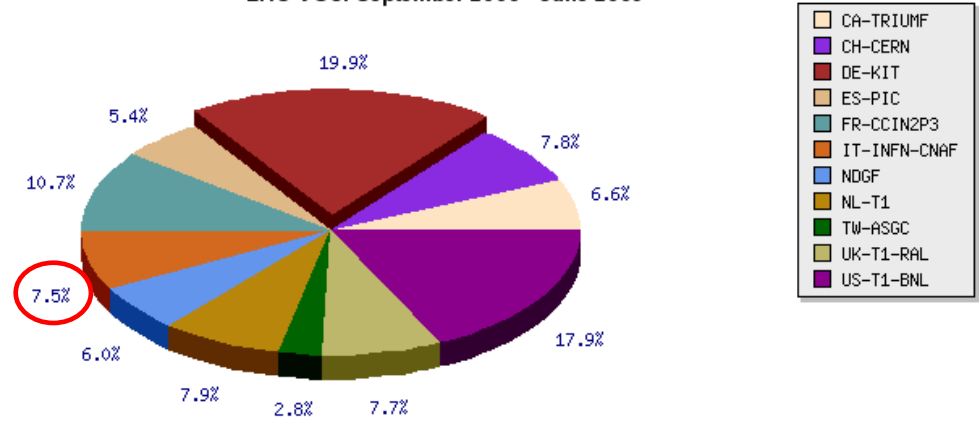


**COUNTRY Normalised Elapsed time per COUNTRY**  
LHC VOs. September 2008 - June 2009



- Canada
- Denmark
- France
- Germany
- Italy
- Netherlands
- Poland
- Romania
- Russia
- Slovenia
- Spain
- Switzerland
- United Kingdom
- United States of America
- Others

**TIER1 Normalised Elapsed time per TIER1**  
LHC VOs. September 2008 - June 2009

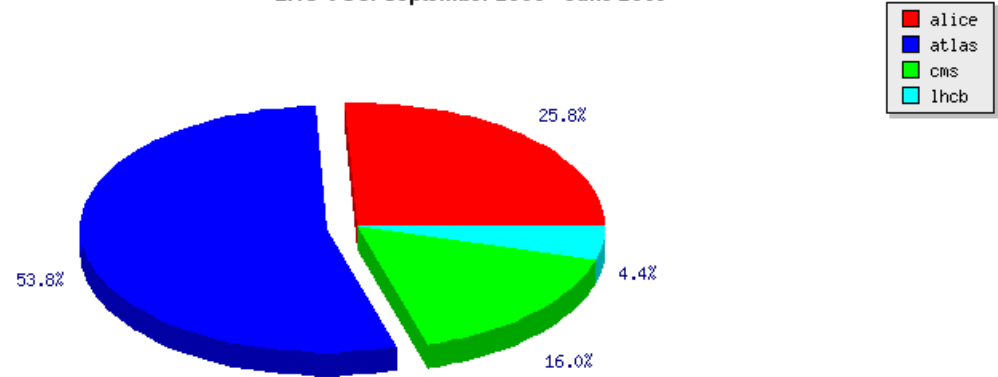


- CA-TRIUMF
- CH-CERN
- DE-KIT
- ES-PIC
- FR-CCIN2P3
- IT-INFN-CNAF
- NDGF
- NL-T1
- TW-ASGC
- UK-T1-RAL
- US-T1-BNL

EE View: TIER1 / normlap / 2008-9-2009:6 / VO-TIER1 / lhc (x) / ACCBAR-LIN / i

2009-06-11 10:12 UTC

**IT-INFN-CNAF Normalised Elapsed time per VO**  
LHC VOs. September 2008 - June 2009



- alice
- atlas
- cms
- lhcb

EE View: IT-INFN-CNAF / normlap / 2008-9-2009:6 / VO-SITE / lhc (x) / ACCBAR-LIN / i

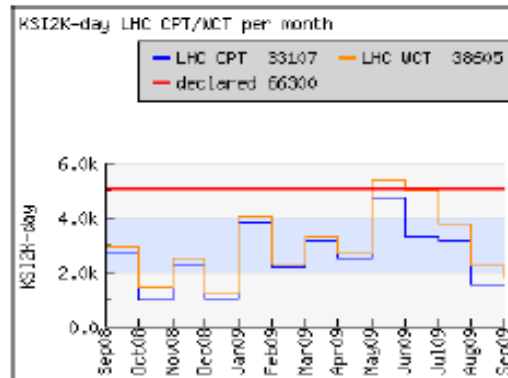
2009-06-11 10:12 UTC

# Il Portale di Monitoring

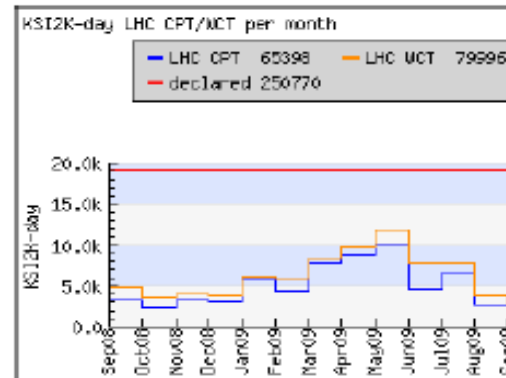


## Uso CPU T2 ATLAS

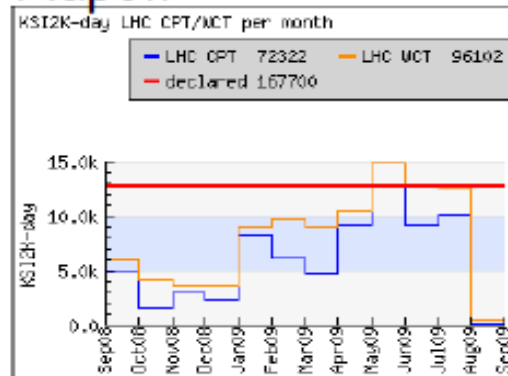
### "Frascati"



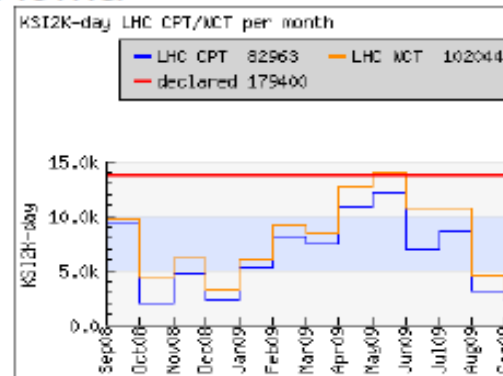
### Milano



### Napoli



### Roma



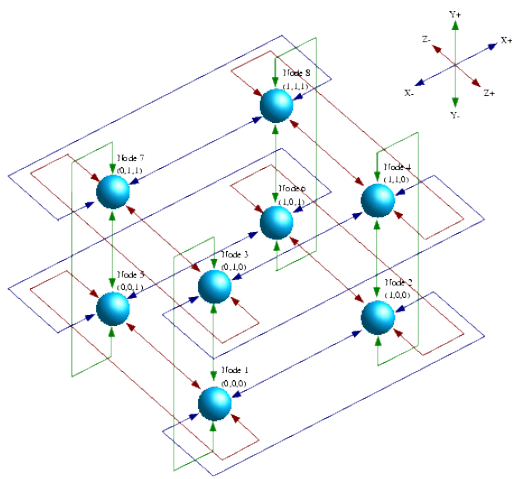
# Conclusioni (1)

Tutto quanto e' sviluppato e gira su Scientific Linux Cern  
(SLC 4 / SLC 5)

E' open source

Puo' essere installato sul proprio laptop o desktop

(ma non e' che giri proprio al primo colpo ...)



# Dalla Fisica Teorica (!?) al Super-Computing ovvero il progetto APE



N. Cabibbo

Progetto INFN, ora collaborazione con  
DESY Zeuthen e Université Paris-Sud 11

	APE	APE100	APEmille	APEnext
Year	1984-1988	1989-1993	1994-1999	2000-2005
Number of processors	16	2048	2048	4096
Topology	Flexible 1D	Next Neighbour 3D	Flexible 3D	Flexible 3D
Total Memory	256 MB	8 GB	64 GB	1 TB
Clock	8 MHz	25 MHz	66 MHz	200 Mhz
Peak Processing Power	1 GFlops	100 GFlops	1 TFlops	7 TFlops

# BACKUP

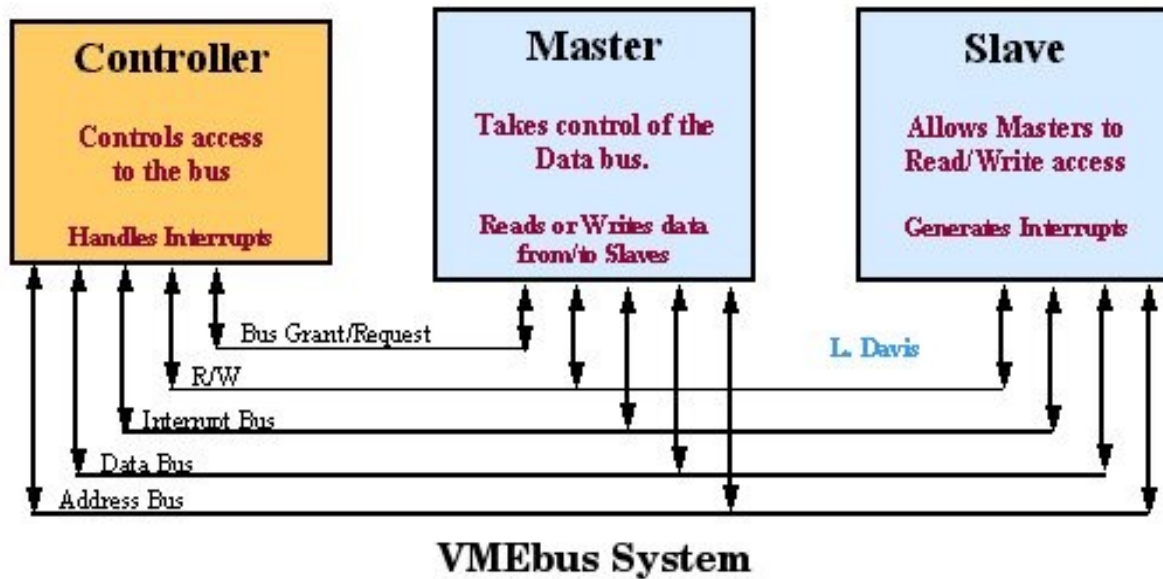
# VME (= Versa Module Eurocard)

- standard (aperto) definito nel 1981
- supporto per multiprocessing in parallelo
- crate modulari di 21 slot interconnesse
- standard elettrico TTL:  $V(0)=+5V$ ,  $V(1)=0V$
- bus dati e di indirizzi a 32 bit
- portata: 40 - 80 (- 160 - 500) MB/s

mappato sul processore 68000

implementazione del bus VERSAbus (definito nel 1979, dalla Motorola, per il processore 68000) su stampati formato Eurocard (100x160 mm)

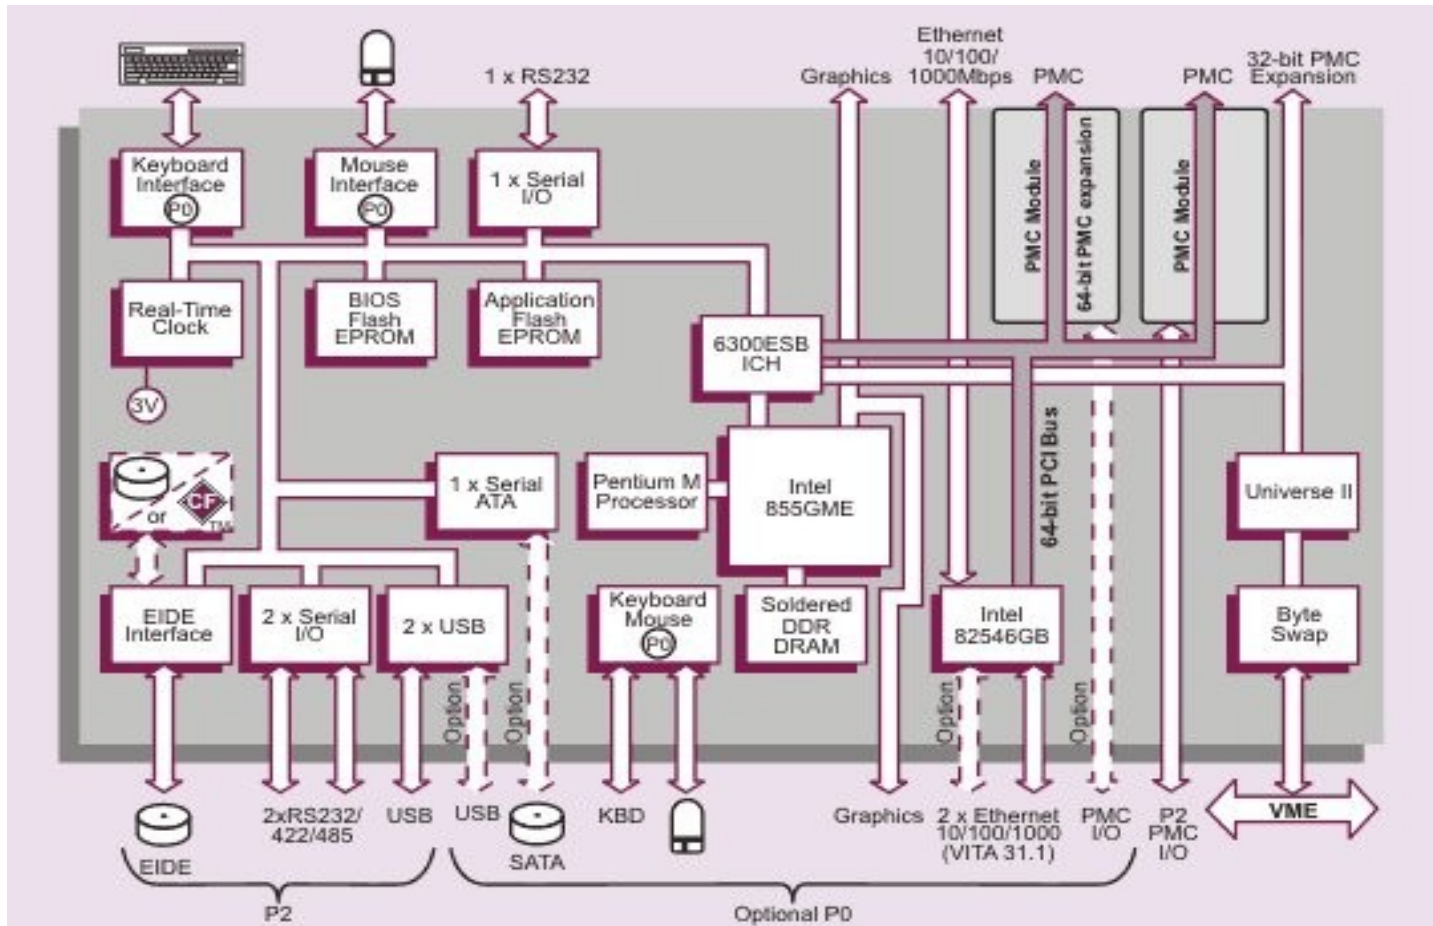
# VME bus



- bus dati: 8-16-24-32-64 bit
- bus indirizzi: 16-24-32-40-64 bit
- 7 livelli di interrupt
- disegnato sul microprocessore Motorola 68000

# S.B.C.

(CCT-VP315/VP317)





# Comunicazioni ...

"Message Passing Interface":

protocollo per generazione e gestione messaggi in rete

"Common Object Request Broker Architecture":

scambio di oggetti indipendentemente da linguaggio e distribuzione nella rete dei soggetti ("broker": mediatore)

(MPI ha un set di modi di comunicazione piu' variegato)

# Multi-threading

suddivisione di un processo in piu' branche in esecuzione parallela (thread)

→ intercomunicazione piu' facile

→ interferenze distruttive piu' facili

sincronizzazione, scheduling (POSIX pthread):

→ create, join, exit

→ mutex (lock, unlock, trylock, ...)

→ condition (wait, signal, ...)

# Multi-core CPU

(un core == una unita' di processamento)

velocita' CPU limitata dal dissipamento:

$$\text{potenza} = a * (\text{freq clock})^2$$

scappatoia(!): + core su di un singolo chip

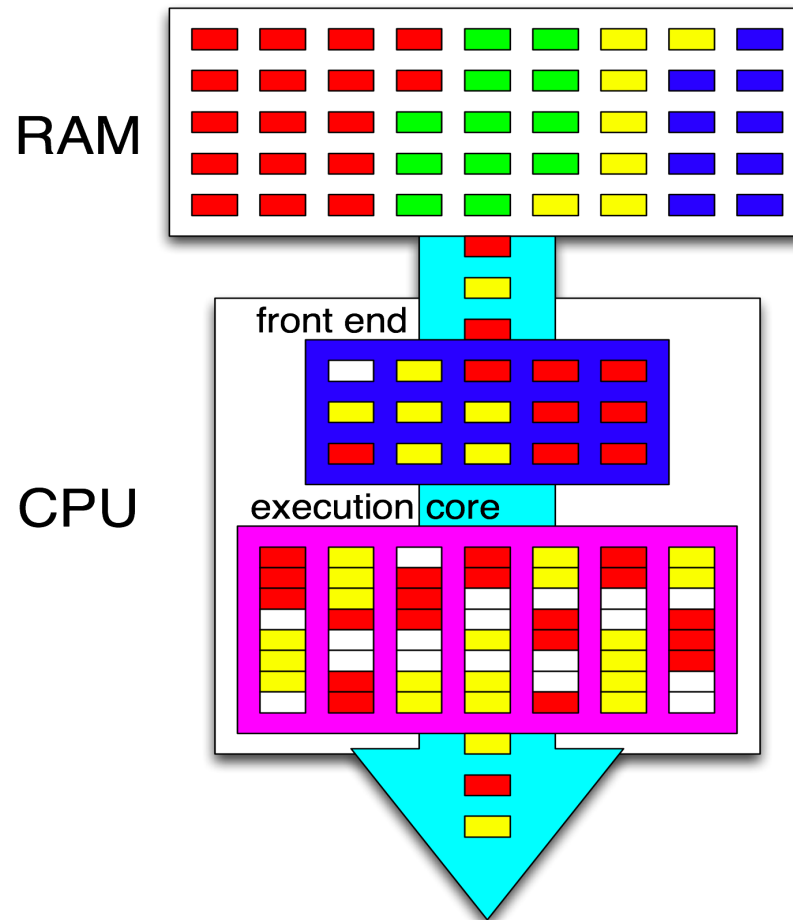
2-3-4-6-8-...

e + thread su un singolo core (tipicamente 2)

1 CPU : 2-4-8-16 thread in parallelo

# HyperThreading CPU

1 CPU : 2-4-8-16  
thread in parallelo

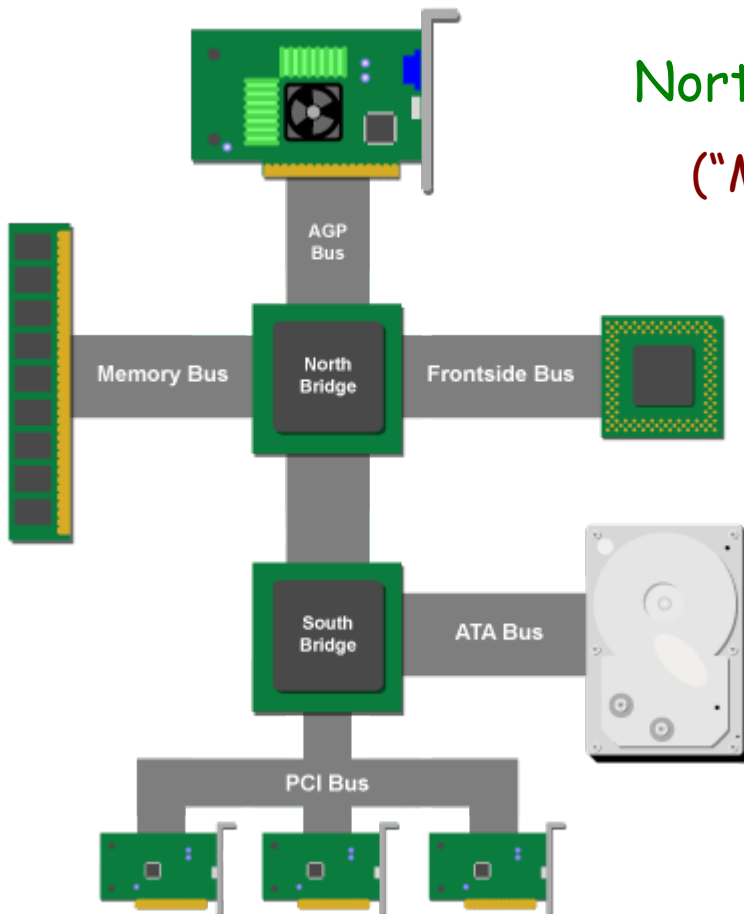


# PC Architecture (intel)

NorthBridge (piu' veloce):

("Memory Controller Hub" - MCH)

RAM, AGP (video)



SouthBridge (piu' lento):

("I/O Controller Hub 2" - ICH2)

altri device (tastiera, mouse, ...)

# Architetture CPU

## CISC (Complex Instruction Set Computer):

- set di istruzioni del microprocessore ampio
- codice assembler piu' vicino al codice di alto livello
- programmi piu' compatti
- istruzioni lente, accessi ripetuti alla memoria

.... ma nel 90% del tempo la CPU utilizza sempre un ristretto sottoinsieme di istruzioni →

## RISC (Reduced Instruction Set Computer):

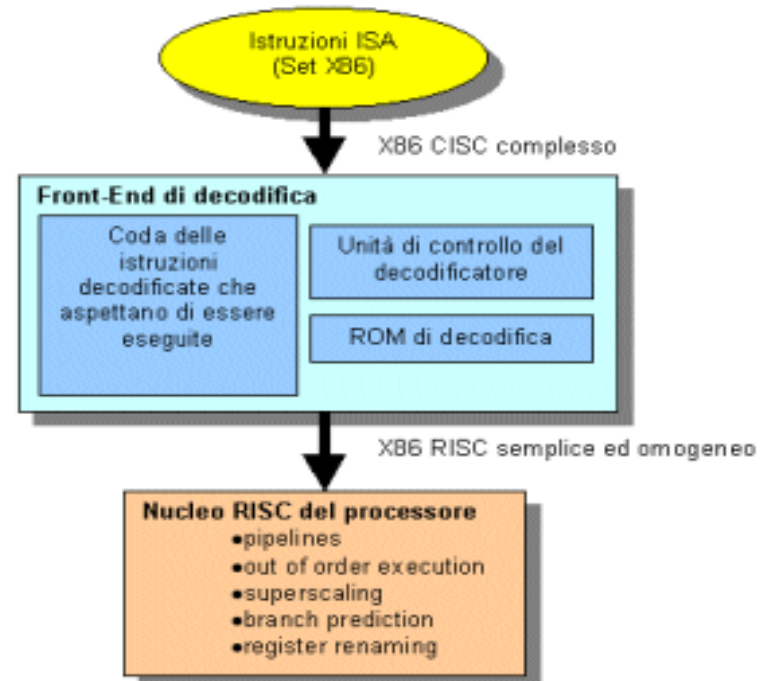
- set limitato di istruzioni semplici e veloci
- cache locale (di primo e secondo livello) per dati e istruzioni
- pipeline (esecuzione parallela di piu' istruzioni)

# Architetture CPU (2)

CISC: VAX, 68000, 80x86

RISC: alpha, sparc, mips, powerpc, arm, pa-risc

Pentium4:  
finge di essere un CISC  
(compatibilita' x86)  
ma lavora come un RISC

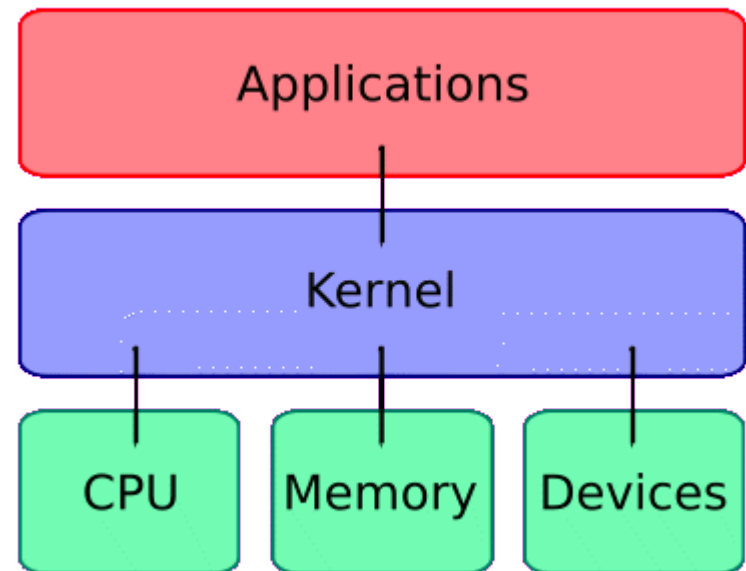


# Kernel

Cuore del sistema operativo

Gestione risorse hardware

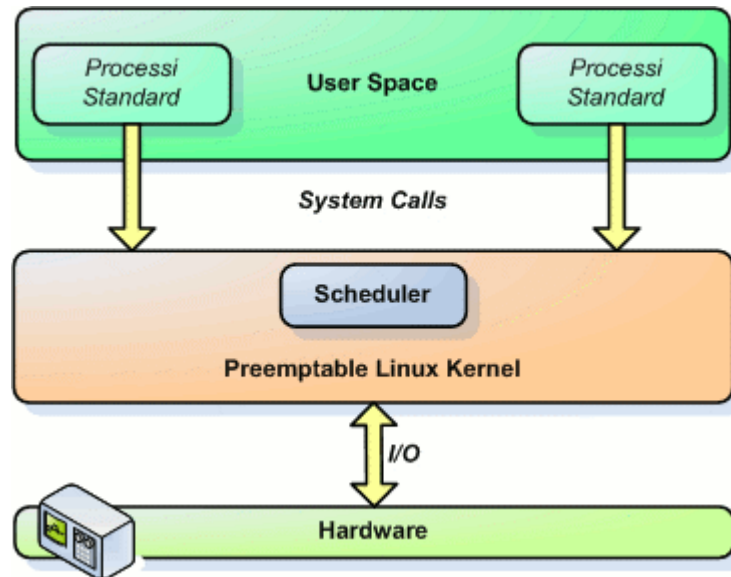
Scheduling e gestione  
processi sw





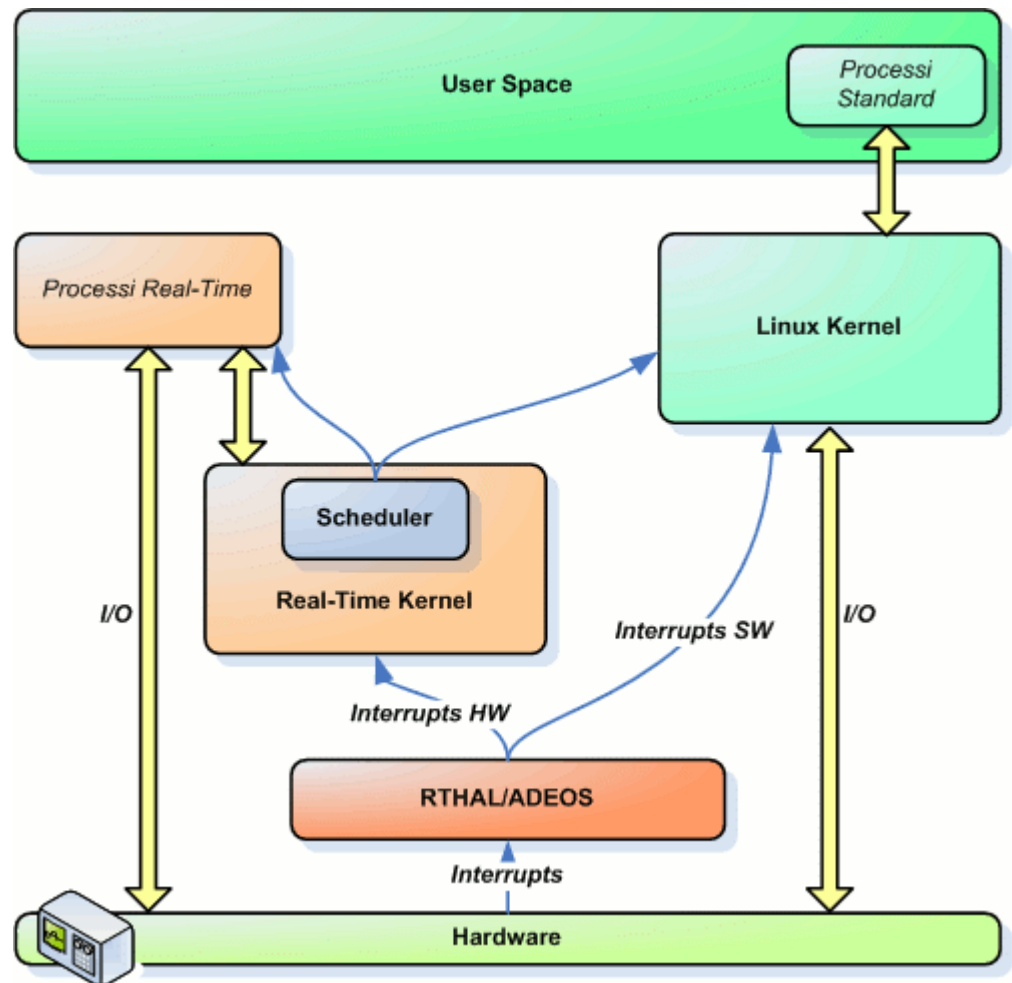
# Real Time (1)

Low-latency patch  
(Ubuntu Studio):  
linux kernel  
interrompibile



# Real Time (2)

**RTAI:** il kernel linux gira come una applicazione a prioritata' piu' bassa



# DAQ Panel

DAQPanel (on pc-atlas-cr-12.cern.ch)

Insert Here Some Info

Setup Script: /sw/tdaq/setup/setup\_tdaq-02-00-03.sh

Part Name: ATLAS

Database File: /atlas/oks/tdaq-02-00-03/combined/partitions/ATLAS.data.xr

Setup Opt:

Oks Opt:

MRS Filter:

OHP Opt: -c

BUSY Opt:

OMD Opt:

TriP Opt: -c

Get Default, Read Info, Get Partition

Main | Mon Advanced | Ctrl Advanced

Start Partition, Monitor Partition, RC Status, Local Procs

OKS, DVS, Log Manager, MRS

Busy, DQM Display, Trigger Presenter, Event Dump

OHP, OMD, ISPY, SFO Display

Log Messages

```
oks_data_editor options --->
ohp options ---> -c
TriP options ---> -c
BUSY options --->
OMD options --->
MRS filter --->
```

Resize, Clear Log, Change role, Exit

You are robertof and your role is

# Run Control

## Macchina a Stati Finiti

**RELOAD CONFIGURATION**

**START/STOP FLOW**

**PROCESSES RUNNING.** Should coincide with the RUN CONTROL STATE

**ERROR LOGGER.** Messages for experts so far...

ATLAS TDAQ Software Graphical User Interface - Expert Control

File Commands Access Control Tools Settings Help

Partition SCT2

Run control

RUN CONTROL STATE: **RUNNING**

Shutdown Boot

Unloading Config

Stop Set

Pause Continue

Run Information

Run type: Calibration

Recording: Disable

Run Start Time: 12/02/08 16:07:07

Run Stop Time:

Total run time: 00:05:43

	Number	Rate
Level 1	0	0 MHz
Level 2	0	0 MHz
Event Builder	0	0 MHz
Event Filter	0	0 MHz

Segment & Resource | Data Set Tags | SCT Supervisor | Infrastructure

Run Control | Run Parameter | MRS | PMG | DataFlow | Monitor

**RUNNING** RootController

- RUNNING** SCT
  - UP** ArchivingServiceImpl
  - UP** SctApiDDCServer
  - UP** SctApiServer
  - RUNNING** SCTConfigurationSen
  - RUNNING** CalibrationController
  - RUNNING** SctEndcapASegmentC
  - RUNNING** ddc\_sct\_ct1

RootController

APPLICATION STATUS: **UP**

RUN CONTROL STATE: **RUNNING**

BUSY STATUS: **FREE**

FAULT STATUS: **OK**

COMMAND: PUBLISH

Operational Monitoring

publish state | publish statistics

Debug Level Control

0 | Set debug level

Membership: IN

Recovery commands: retry, ignore error, clear error

Status | Informations | Commands

16:12:55	WARNING	ArchivingService_...	Sct:IsException	ISException Error writing to IS server. Couldn't write: RetrievedData.SctData:RawScanResult.38121.0.2022013000033410Exception: Error writing to IS server. Couldn't write: RetrievedData.SctData:RawScanResult.38121.0.20220130000334 -- 39 similar messages suppressed, last occurrence was at 2008-Feb-12 16:12:54
16:12:44	WARNING	ArchivingService_...	Sct:IsException	ISException Error writing to IS server. Couldn't write: RetrievedData.SctData:RawScanResult.38121.0.2022013000018810Exception: Error writing to IS server. Couldn't write:

# Running ...

ATLAS TDAQ SOFTWARE - Partition ATLAS

File Commands Access Control Settings Logging Level Help

Commit & Reload

**RUN CONTROL STATE** **RUNNING**

Run Control Commands

SHUTDOWN BOOT  
 TERMINATE INITIALIZE  
 UNCONFIG CONFIG  
 STOP START  
 HOLD TRG RESUME TRG

Run Information & Settings

Lumi Block 258

	Number	Rate
Level 1	343214673	14.16 kHz
Level 2	1524053	121.01 Hz
Event builder	1524050	120.71 Hz
Event filter	1500500	139.31 Hz
Recorded	1937034	151.28 Hz

Information Counters Settings

TGC Status ZDC Trigger Tile LAr DFPanel PMG L1Calo BCM

Run Control Segments & Resources Dataset Tags

- RUNNING** RootController
  - RUNNING** TDAQ:pc-tdq-onl-16
  - RUNNING** GlobalMonitoringSegment:pc-tdq-mon-16
  - RUNNING** RPC
  - RUNNING** Tile
  - RUNNING** LArg
  - RUNNING** TRT
  - RUNNING** **DQMController**
  - RUNNING** MDT
  - RUNNING** Pixel
  - RUNNING** SCT
  - RUNNING** TGC
  - RUNNING** BCM
  - RUNNING** IDG-MonitoringSegn
  - RUNNING** LUCID
  - RUNNING** ZDC

Trigger Rate Presenter

File Plots Help

Plots Tables

L1 + L2 + EF

Rate (Hz) (x1e3)

Time

Infrastructure Advanced

Show Online Segment Find: Match Case Repeats

Subscription criteria WARNING ERROR FATAL INFORMATION Expression Subscribe

TIME	SEVERITY	APPLICATION	NAME	MESSAGE
17:30:45	WARNING	ROS-TIL-LBC-01	ROS::CoreException	Timeout: in request for fragment with L1 ID 3321894837 -- 3482 similar messages suppressed, last occurrence was at 2009-Oct-15 17:30:41
17:30:44	ERROR	CheckBCIDGnam	bcidcheck:AnyError	Run 135161 Ev 179472 Ref 12015 L1 0xc401e45a TT 0x81 BC 0xa0d Status 0x 1 - SCT_BARREL_A_SIDE new partition found -- 478 similar messages suppressed, last occurrence was at 2009-Oct-15 17:30:36
17:30:42	WARNING	TRT	OnlRec:HardwareRec...	"Resync finished. Restarting trigger"
17:30:40	ERROR	TRTSyncController	rc:HardwareSynchron...	TRT would like to resynchronize its hardware.

Clear Message format Number of visible rows 100 Current MRS subscription WARNING|ERROR|FATAL

2009-10-15 15:04  
 next runs:  
 LAR RODC\_EMBA2 to enable back  
 TRT ROS to enable back

# Message Reporting System

MRS Monitor [ATLAS]

Partition: ATLAS

Subscription criteria:  WARNING  ERROR  FATAL  INFORMATION  Expression Subscribe

TIME	SEVERITY	APPLICATION	NAME	MESSAGE
15:39:41	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3599./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:41
15:39:40	WARNING	LVL2-L2-2-rack-...	gatherer::Issue	Histogram'L2PU-5920./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3752./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning
15:39:40	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3176./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:40
15:39:40	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3560./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning
15:39:40	ERROR	CheckBCIDGnam	bcidcheck::AnyError	Run 136207 Evt 18434 Ref 1 L1 0xac003643 TT 0xc0 BC 0x576 Status 0x 1 full 0 - event format error -- 95 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:40	ERROR	CheckBCIDGnam	bcidcheck::AnyError	Run 136207 Evt 18435 Ref 1 L1 0xab00c357 TT 0xc0 BC 0xae Status 0x 1 full 0 - FORWARD_BCM, module=12 (opt=0) (ROB 0x81000c) BCID internal mismatch: 0xbb1 / 0xae -- 95 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3690./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning
15:39:40	WARNING	ROS-TIL-LBC-01	ROS::CoreException	Timeout: in request for fragment with L1 ID 2919260819 -- 2252 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3194./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3199./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3162./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning
15:39:39	WARNING	LVL2-L2-1-rack-...	gatherer::Issue	Histogram'L2PU-8103./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-2-rack-...	gatherer::Issue	Histogram'L2PU-6107./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTHEC' can not be sumed because histograms have incompatible binning -- 10 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:39
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3208./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning
15:39:39	WARNING	SFI-32	SFI::DataIntegrity...	Problem with data integrity: Event fragment from ROB 0x81000c (ROS: BCM_ROS , SubDet: 129) with LVL1ID: 0xae004c0e and BCID: 1357 has a BCID mismatch: Event_BCID = -195. [ ROS Fragment status()= 0x1 ] -- 57 similar messages suppressed, last occurrence was at 2009-Oct-23 15:39:38
15:39:38	WARNING	LVL2-L2-2-rack-...	gatherer::Issue	Histogram'L2PU-6307./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning
15:39:39	WARNING	LVL2-L2-4-rack-...	gatherer::Issue	Histogram'L2PU-3119./EXPERT/CosmicLArCalib_V2LArL2ROBListWriter/RobldTTTEM' can not be sumed because histograms have incompatible binning

Clear  Message format    Number of visible rows: 2,000 Current MRS subscription: WARNING|ERROR|FATAL

# Diario di Bordo ...

ATLOG ATLAS - Konqueror (on pc-atlas-cr-02.cern.ch)

Location Edit View Go Bookmarks Tools Settings Window Help

Location: <https://pc-atlas-www.cern.ch/elog/ATLAS/ATLAS/>

ATLAS  
Electronic logbook for the ATLAS experiment, Page 1 of 2669  
Logged in as "Ferrari Roberto" **ELOG**

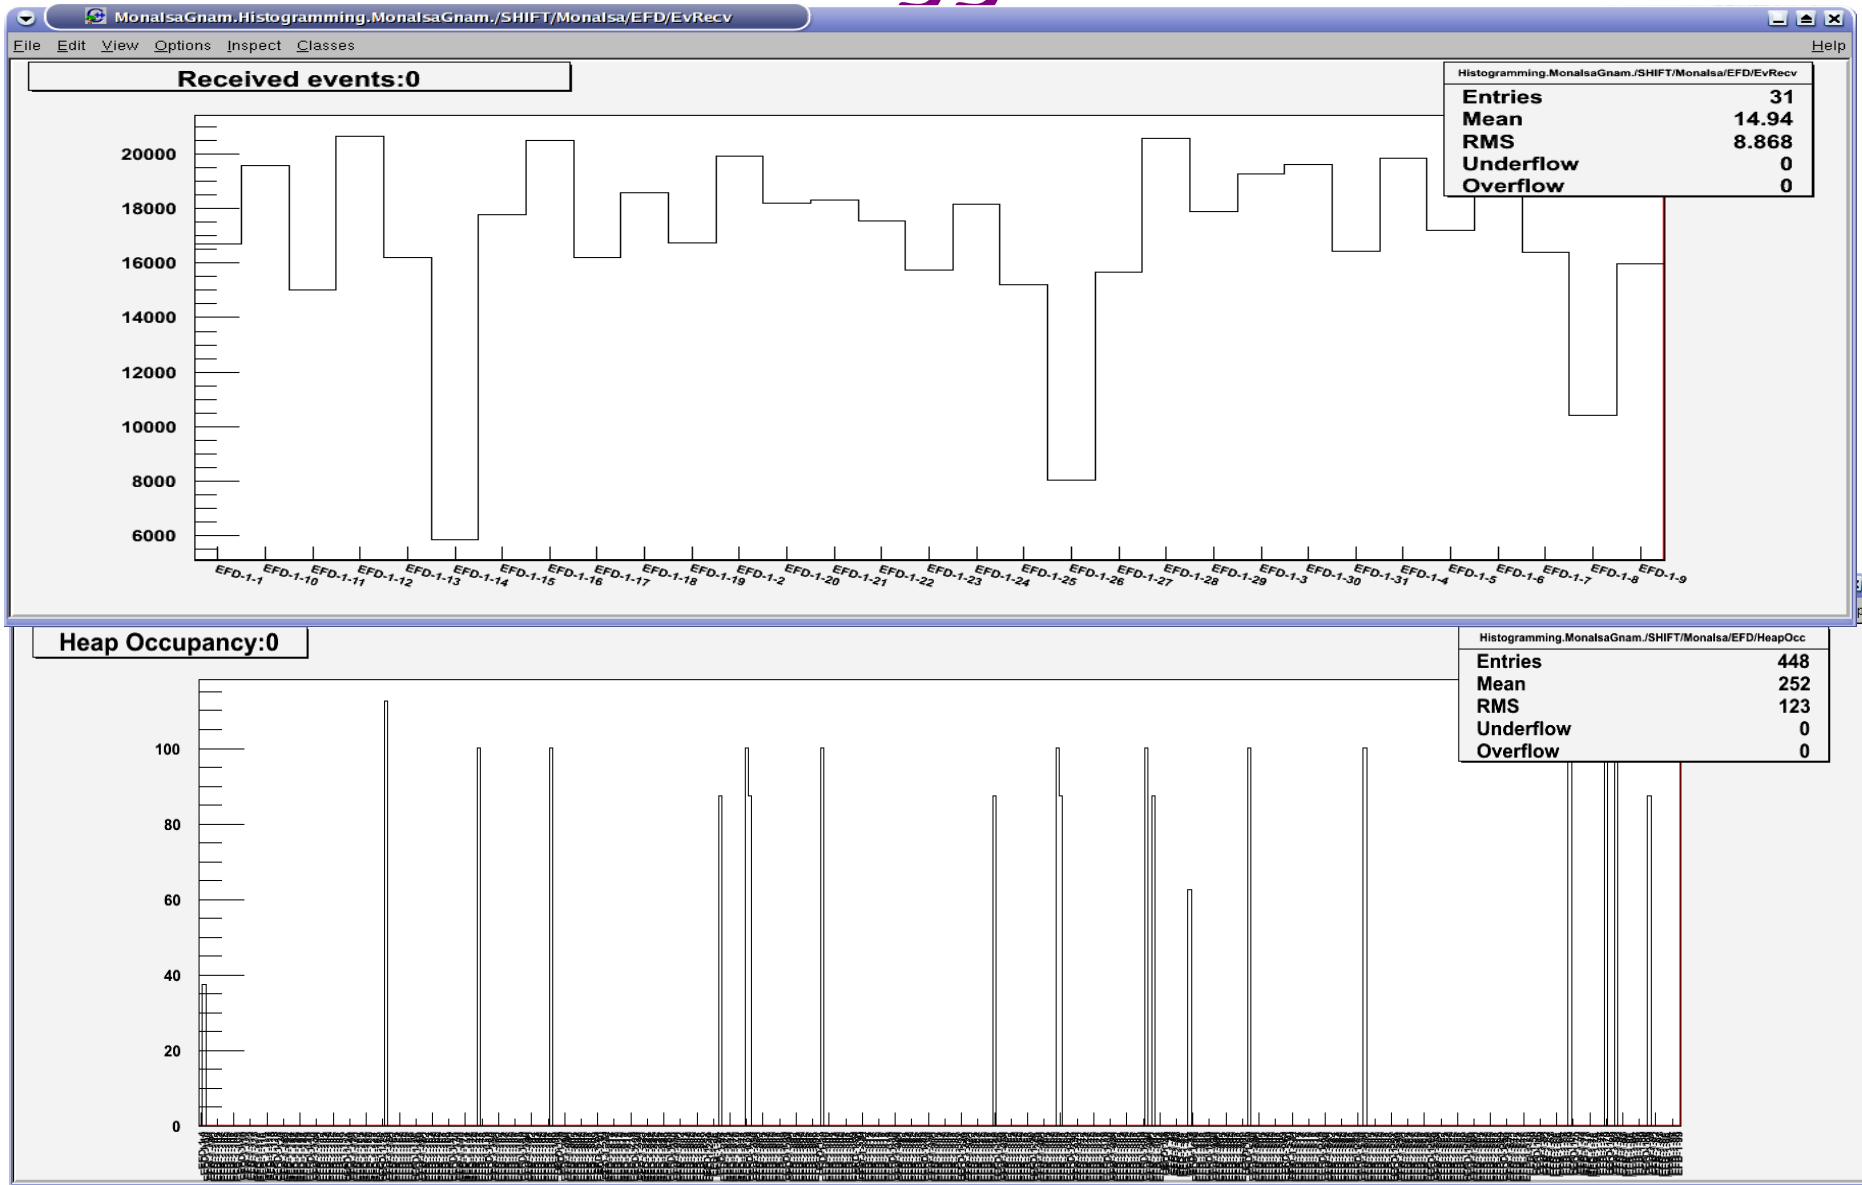
List | Find | New | Select | Logout | Help

Full | Summary | Threaded  All entries Message Type: All entries System Affected: All entries **53375 Entries**

Goto page 1, 2, 3 ... 2667, 2668, 2669 Next

Date	Author	Message Type	System Affected	Subject	Text
23.10.09 15:29	Canepa Anadi	Data Quality		Online DQ Shifter Summary	Summary of Online DQ Shifter Atlantis and VPl running fine
23.10.09 15:26	DCS-IS	Slimos-TI	DSS   Tech. Infra	MAG_Toroid_SlowDump	DSS Alarm: MAG_Toroid_SlowDump
23.10.09 15:17	Tikhomirov Vladimir.O	Default Message Type	TRT	New shift	Combined cosmic run #136207 is going lh15m. TRT status OK.
23.10.09 15:14	Bondioli Bond	Shift Summary	DAQ	Shift Summary for Run Control desk	Got run 136183 from previous shift - GM-Application are constantly dying
23.10.09 15:11	Hayakawa Takashi	Default Message Type	TGC	TGC on-call shift report	*** Readout and Trigger timing *** run#136176 run#136183
23.10.09 15:10	Perez Reale Valeria	Shift Summary	TRT   Tile   Cryo   DCS   Pixel   Magnets   TGC   ID Gen. (IC)   MDT   DSS   SCT   LArg   Network   RPC   CSC   DAQ   HLT   LVL1   Monitoring   GAS   SysAdmins   BCM   Lucid   Counting Room   Tier0   Beam Conditions   ZDC   Event Displays	Shift Summary for Shift Leader	07:00 Run from started at 4am last night is still running: [B] Run 136183 (620, 881, 775 ) [/B] . Initially,
23.10.09 15:06	Kim Tae Jeong	Shift Summary	HLT   LVL1	Shift Summary for Trigger desk	Run upon arrival: #136183 Going since 04:00
23.10.09 15:05	Qi Ming	Shift Summary	TGC	Shift Summary for Muon Desk 3 - TGC	Shift summary, 07:00-15:00, 23 Oct. 2009 Run136183, Begin at 03:55:26, Run continued,
23.10.09 15:04	Dubbert Joerg	Default Message Type	DCS   MDT	BOG6A12 ML1 HV interlock asserted	Asserted HV interlock for BOG6A12 ML1. Chamber tripped repeatedly since yesterday. Stable for some time (up to some hours) then fast
23.10.09 15:03	Ghodbane Nabil	Tile	Tile	end of shift summary	Ongoing run: 136207 List of good runs with TileCal in : Summary of what happened during the shift
23.10.09 15:02	Ferretti Claudio	Shift Summary	TGC   MDT   RPC   CSC	Shift Summary for Muon Desk 1 - MDT/CSC	Many runs (stooped mostly because of RC problems) 136183 ~ 160 Kev. Run 136183 ended at about 9:00. Sectors in

# Monitoraggio Online





# Information Service

Partition 'ATLAS', server 'DF-EF-Segment-01-rack-Y03-06D2-iss'

Name	Type	Modified	Description
EFD-1-25	EFD	16/7/08 09:43:31.549965	
EFD-1-26	EFD	16/7/08 09:43:34.503773	
EFD-1-27	EFD	16/7/08 09:43:31.834124	
EFD-1-28	EFD	16/7/08 09:43:31.946579	

Value	Type	Name	Description
pc-tdq-xpu-0245:/local_L/efHeap/sharedHeap.crrc.ATLAS	String	SharedHeap	SharedHeap file fullpath
3	U16	ConnNbrSFIs	Number of connected SFIs (sum over InputTasks)
5	U16	ConnNbrSFOs	Number of connected SFOs (sum over OutputTasks)
4	U16	ConnNbrPTs	Number of connected PTs (sum over ExtPTsTasks)
87.54	Double	HeapOcc	SharedHeap occupancy (%)
1521	S32	EventsRcv	Number of received events
1514	S32	EventsSent	Number of events sent to SFO (ie: Dismissed-Deleted)
7	S32	EventsInside	Number of events inside
0	U32	EventsWaitingForProc	Number of events waiting for processing
2	U32	EventsWaitingForDeli	Number of events waiting to be sent to SFO
0	Double	RateIn	Current rate of incoming events (Hz)
0	Double	RateOut	Current rate of events sent to SFO (Hz)
0	Double	FluxIn	Current rate of space allocation in SH; >~ data flu
0	Double	FluxOut	Current data flux to SFO (MB/s)
-1	Double	FlowCtrlIStopTime	Guess of the stop transition time (s)
460	U32	FlowCtrlISleepTime	Current flow control sleep time (ms)
538	U32	FlowCtrlBarrierLocks	Number of times the input barrier has been locked
0	S32	ptioNbrProcTimeouts	Number of processing timeouts
0	S32	ptioNbrSocketHungUps	Number of PT socket hung-ups
0	S32	ptioNbrForceAccept	Number of force accepted events
0	S32	efioNbrSfiBrokenConn	Number of broken connections to SFI
0	S32	efioNbrSfoBrokenConn	Number of broken connections to SFO
1521, 0, 0, 0, 0, 0	S32[6]	EventTagTypesIn	Type counters: phys, calib, reserved, debug, unknow
1519, 0, 0, 2, 0, 0	S32[6]	EventTagTypesOut	Type counters: phys, calib, reserved, debug, unknow

2

403 objects | 24 attributes

# NAGIOS

**Nagios®**

**General**

- Home
- Documentation

**Current Status**

- Tactical Overview
- Map
- Hosts
- Services
- Host Groups
  - Summary
  - Grid
- Service Groups
  - Summary
  - Grid
- Problems
  - Services (Unhandled)
  - Hosts (Unhandled)
  - Network Outages

Quick Search:

**Reports**

- Availability
- Trends
- Alerts
  - History
  - Summary
  - Histogram
- Notifications
- Event Log

**System**

- Comments
- Downtime
- Process Info
- Performance Info

**Current Network Status**  
 Last Updated: Tue Oct 27 14:35:17 CET 2009  
 Updated every 90 seconds  
 Nagios® Core™ 3.2.0 - [www.nagios.org](http://www.nagios.org)  
 Logged in as *atruser*

[View History For all hosts](#)  
[View Notifications For All Hosts](#)  
[View Host Status Detail For All Hosts](#)

**Host Status Totals**

Up	Down	Unreachable	Pending
35	0	0	0
<b>All Problems</b>		<b>All Types</b>	
0		35	

**Service Status Totals**

Ok	Warning	Unknown	Critical	Pending
648	14	0	5	0
<b>All Problems</b>		<b>All Types</b>		
19		667		

## Service Status Details For All Hosts

Host	Service	Status	Last Check	Duration	Attempt	Status Information
pcatb120	/scratch	OK	10-27-2009 14:34:30	6d 2h 10m 2s	1/3	DISK OK - free space: /scratch 1764 MB (97% inode=99%):
	atmp	OK	10-27-2009 14:34:30	6d 2h 10m 2s	1/3	DISK OK - free space: atmp 1764 MB (93% inode=99%):
	AFS	OK	10-27-2009 14:33:52	6d 1h 29m 58s	1/3	AFS OK: AFS mounted and daemon available
	Automount	OK	10-27-2009 14:34:26	6d 2h 1m 59s	1/3	PROCS OK: 1 process with command name 'automount'
	CFEngine	OK	10-27-2009 14:33:52	6d 1h 29m 58s	1/3	CFENGINE OK: Latest run was made only 0.6 hours ago
	Current Load	OK	10-27-2009 14:33:53	2d 0h 57m 44s	1/3	OK - load average: 0.02, 0.01, 0.00
	Current Users	OK	10-27-2009 14:34:30	5d 1h 24m 5s	1/3	USERS OK - 2 users currently logged in
	Memory	OK	10-27-2009 14:34:34	2d 0h 56m 9s	1/3	MEMORY OK - 8712 kb (1.7%) free, 145328 kb (28.2%) cached, 154040 kb (29.9%) available
	NFS Server	OK	10-27-2009 14:34:28	2d 0h 59m 8s	1/3	PROCS OK: 8 processes with command name 'nfsd'
	NTP Daemon	OK	10-27-2009 14:34:28	2d 0h 59m 8s	1/3	PROCS OK: 1 process with command name 'ntpd'
	PMG Server	OK	10-27-2009 14:34:28	2d 0h 59m 8s	1/3	PROCS OK: 0 processes with command name 'pmgserver'
	Ping	OK	10-27-2009 14:33:53	22d 5h 41m 40s	1/3	PING OK - Packet loss = 0%, RTA = 0.14 ms
	SSH	OK	10-27-2009 14:34:30	6d 1h 28m 16s	1/3	SSH OK - OpenSSH_4.3 (protocol 2.0)
	TCP	OK	10-27-2009 14:33:51	6d 1h 29m 58s	1/3	TCP OK: TX errors at 0.0% RX errors at 0.0%
	Time	OK	10-27-2009 14:34:32	6d 2h 24m 0s	1/3	TIME OK - 0 second time difference
Total Processes	OK	10-27-2009 14:34:32	2d 0h 58m 35s	1/3	PROCS OK: 142 processes	
UDP	OK	10-27-2009 14:33:53	4d 11h 24m 17s	1/3	UDP OK: Receive errors at 0.0%	
Zombies	OK	10-27-2009 14:34:26	2d 0h 59m 8s	1/3	PROCS OK: 1 process with STATE = Z	
root disk	OK	10-27-2009 14:34:30	6d 2h 24m 2s	1/3	DISK OK - free space: / 4938 MB (52% inode=90%):	
pcatb121	/scratch	OK	10-27-2009 14:34:26	22d 5h 41m 12s	1/3	DISK OK - free space: /scratch 1306 MB (97% inode=99%):
	atmp	OK	10-27-2009 14:34:29	15d 6h 5m 38s	1/3	DISK OK - free space: atmp 1844 MB (98% inode=99%):
	AFS	OK	10-27-2009 14:33:52	15d 6h 5m 37s	1/3	AFS OK: AFS mounted and daemon available
	Automount	OK	10-27-2009 14:34:34	15d 6h 5m 36s	1/3	PROCS OK: 1 process with command name 'automount'
	CFEngine	OK	10-27-2009 14:33:52	15d 5h 32m 39s	1/3	CFENGINE OK: Latest run was made only 0.6 hours ago

# www.nagios.org

**NAGIOS DISTRIBUTION**

