

CHEP04

Monday 27 September 2004 - Friday 01 October 2004

Interlaken, Switzerland



Recueil des résumés

This document and all information concerning the conference programme can be accessed online via the CHEP04 web site: <http://www.chep04.org>. The information on the web will be updated in the event on any last minute changes in the programme. These will also be communicated directly at the conference.

This document contains all abstracts of the talks and posters to be presented at CHEP 2004. They are ordered by Abstract Number. Information on when, and in which session, the presentation is due to be made is also given here for completeness. The Conference Programme should be consulted to view the complete schedule of talks and cross reference to this document made using the Abstract Number.

Contents

”RecPack”, a general reconstruction tool-kit	1
50 years of Computing at CERN	1
64-Bit Opteron systems in High Energy and Astroparticle Physics	2
A Control Software for the ALICE High Level Trigger	2
A Data Grid for the Analysis of Data from the Belle Experiment	3
A Dynamically Reconfigurable Data Stream Processing System	3
A GRID approach for Gravitational Waves Signal Analysis with a Multi-Standard Farm Prototype	4
A Gaussian-sum filter for vertex reconstruction	4
A Globally Distributed Real Time Infrastructure for World Wide Collaborations	5
A Hardware Based Cluster Control And Management System	6
A Level-2 trigger algorithm for the identification of muons in the Atlas Muon Spectrometer	6
A Lightweight Monitoring and Accounting System for LHCb DC’04 Production	7
A Multidimensional Approach to the Analysis of Grid Monitoring Data	7
A New STAR Event Reconstruction Chain	8
A Pattern-based Continuous Integration Framework for Distributed EGEE Grid Middle- ware Development	8
A Regional Analysis Center at the University of Florida	9
A Scalable Grid User Management System for Large Virtual Organization	10
A database perspective on CMS detector data	10
A database prototype for managing computer systems configurations	11
A distributed, Grid-based analysis system for the MAGIC telescope	11
A general and flexible framework for virtual organization application tests in a grid system	12

A high-level language for specifying detector coincidences	12
A kinematic and a decay chain reconstruction library	13
A statistical toolkit for data analysis	13
AIDA, JAIDA and AIDAJNI: Data Analysis using interfaces	14
ALICE Multi-site Data Transfer Tests on a Wide Area Network	15
AMS-02 Computing and Ground Data Handling	16
ARDA Project Status Report	16
ATLAS Data Challenge Production on Grid3	17
ATLAS Detector Description Database Architecture	17
ATLAS Distributed Analysis	18
ATLAS Metadata Interfaces (AMI) and ATLAS Metadata Catalogs	19
ATLAS Production System in ATLAS Data Challenge 2	19
Adaptive Multi-vertex fitting	20
Adding Kaons to the Bertini Cascade Model	21
Addressing the persistency patterns of the time evolving HEP data in the ATLAS/LCG MySQL Conditions Databases	21
AliEn Web Portal	22
Alibaba: A heterogeneous grid-based job submission system used by the BaBar experiment	22
An Embedded Linux System Based on PowerPC	23
An Object-Oriented Simulation Program for CMS	23
An intelligent resource selection system based on neural network for optimal application performance in a grid environment	24
Application of the SAMGrid Test-Harness for Performance Evaluation and Tuning of a Distributed Cluster Implementation of Data Handling Services	25
Applications of the FLUKA Monte Carlo code in High Energy and Accelerator Physics	26
Aspect-Oriented Extensions to HEP Frameworks	27
Aspects	27
Authentication/Security services in the ROOT framework	28
AutoBlocker: A system for detecting and blocking of network scanning based on analysis of netflow data.	28
Automated Tests in NICOS Nightly Control System	29

Automatic Procedures as Generated Analysis Tool	29
BFD: A software management & deployment tool for mixed-language distributed projects	30
BaBar Bookkeeping - a distributed meta-data catalog of the BaBar event store.	31
BaBar computing - From collisions to physics results	32
BaBar simulation production - A millennium of work in under a year	32
Bayesian Approach for Combined Particle Identification in ALICE Experiment at LHC	33
Beyond Persistence: Developments and Directions in ATLAS Data Management	34
Boosting the data logging rates in run 4 of the PHENIX experiment	34
Breaking the 1 GByte/sec Barrier? High speed WAN data transfers for science	35
Bringing High-Performance Networking to HEP Users	35
Building Global HEP Systems on Kerberos	36
Building the LCG: from middleware integration to production quality software	36
CASTOR: Operational issues and new Developments	37
CDB - Distributed Conditions Database of BaBar Experiment	38
CERN Modular Physics Screensaver or Using spare CPU cycles of CERN's Desktop PCs	38
CERN's openlab for Datagrid applications	39
CHIPS based hadronization of quark-gluon strings	39
CHOS, a method for concurrently supporting multiple operating system	40
CLHEP Infrastructure Improvements	40
CMD-3 Project Offline Software Development	41
CMS Detector Description: New Developments	41
CMS Software Installation	42
CMS Tracker Visualisation Tools	42
Carrot ROOT Apache Module	43
Central Reconstruction System on the RHIC Linux Farm in Brookhaven Laboratory	43
Chimera - a new, fast, extensible and Grid enabled namespace service	44
Cluster architectures used to provide CERN central CVS services	45
Comparative study of the power of goodness-of-fit algorithms	45
Composite Framework for CMS Applications	46
Computing Models and Data Challenges of the LHC experiments	46

Computing for Belle	46
Concepts and technologies used in contemporary DAQ systems	47
Conditions Databases: the interfaces between the different ATLAS systems	47
Conference Closeout	48
Conference Conclusions	48
Control in the ATLAS TDAQ System	48
Control and state logging for the PHENIX DAQ System	49
Cross Experiment Workflow Management: The Runjob Project	50
Current Status of Fabric Management at CERN	50
D0 data processing within EDG/LCG	51
DIRAC - The Distributed MC Production and Analysis for LHCb	52
DIRAC Lightweight information and monitoring services using XML-RPC and Instant Mes- saging	53
DIRAC Workload Management System	53
DZERO Data Acquisition Monitoring and History Gathering	54
Data Management in EGEE	54
Data Rereprocessing on Worldwide Distributed Systems	55
Data management services of NorduGrid	56
Database Usage and Performance for the Fermilab Run II Experiments	56
Defining a Semantic Web Initiative for High Energy Physics	57
Deploying and operating LHC Computing Grid 2 (LCG2) During Data Challenges	57
Deployment of SAM for the CDF Experiment	58
Design and Implementation of a Notification Model for Grid Monitoring Events	58
Designing a Useful Email System	59
Detector-independent vertex reconstruction toolkit (VERTIGO)	59
Deterministic Annealing for Vertex Finding at CMS	60
Developing & Managing a large Linux farm - the Brookhaven Experience	60
Development and use of MonALISA high level monitoring services for the star unified Meta-Scheduler	61
Development of algorithms for cluster finding and track reconstruction in the forward muon spectrometer of ALICE experiment	62

Developments of Mathematical Software Libraries for the LHC experiments	62
Disk storage technology for the LHC T0/T1 centre at CERN	63
Distributed Computing Grid Experiences in CMS DC04	63
Distributed Filesystem Evaluation and Deployment at the US-CMS Tier-1 Center	64
Distributed Testing Infrastructure and Processes for the EGEE Grid Middleware	65
Distributed Tracking, Storage, and Re-use of Job State Information on the Grid	66
Distributed computing and oncological radiotherapy: technology transfer from HEP and experience with prototype systems	67
Domain Specific Visual Query Language for HEP analysis or How far can we go with user friendliness?	67
Don Quijote - Data Management for the ATLAS Automatic Production System	68
Dynamic Matched Filters for Gravitational Waves Detection	69
EU Grid Research - Projects and Vision	69
Enterasys - Networks that Know	69
Evaluation of Grid Security Solutions using Common Criteria	70
Event Data Model in ATLAS	70
EventStore: Managing Event Versioning and Data Partitioning using Legacy Data Formats	71
Evolution and Revolution in the Design of Computers Based on Nanoelectronics	71
Evolution of LCG-2 Data Management	72
Experience integrating a General Information System API in LCG Job Management and Monitoring Services	73
Experience producing simulated events for the DZero experiment on the SAM-Grid	73
Experience with CORBA communication middleware in the ATLAS DAQ	74
Experience with Deployment and Operation of the ATLAS Production System and the Grid3+ Infrastructure at Brookhaven National Lab	74
Experience with POOL from the LCG Data Challenges of the three LHC experiments	75
Experience with Real Time Event Reconstruction Farm for Belle Experiment	75
Experience with the Unified Process	76
Experiences Building a Distributed Monitoring System	77
Experiences with Data Indexing services supported by the NorduGrid middleware	77
Experiences with the gLite Grid Middleware	77

Experiment Software Installation experience in LCG-2	78
Extending EGS with SVG for Track Visualization	79
FAMOS, a FAst MOnte-Carlo Simulation for CMS	79
Fast reconstruction of tracks in the inner tracker of the CBM experiment	80
Fast tracking for the ATLAS LVL2 Trigger	80
Fault Tolerance and Fault Adaption for High Performance Large Scale Embedded Systems	81
Federating Grids: LCG meets Canadian HEPGrid	81
File-Metadata Management System for the LHCb Experiment	82
FroNtier: High Performance Database Access Using Standard Web Components in a Scal- able Multi-tier Architecture	82
From Geant 3 to Virtual Monte Carlo: Approach and Experience	83
Full Event Reconstruction in Java	84
Future processors: What is on the horizon for HEP farms?	84
G-PBox: a Policy Framework for Grid Environments	84
GILDA: a Grid for dissemination activities	85
GROSS: an end user tool for carrying out batch analysis of CMS data on the LCG-2 Grid.	86
Geant4 as Simulation Toolkit addressed to interplanetary manned missions studies: re- quired developments and improvements	86
Geant4: status and recent developments	87
Generic logging layer for the distributed computing	87
Genetic Programming and its application to HEP	88
Gfarm v2: A Grid file system that supports high-performance distributed and parallel data computing	88
Global Distributed Parallel Analysis using PROOF and AliEn	89
Global Grid User Support for LCG	89
Globally Distributed User Analysis Computing at CDF	90
Go4 analysis design	91
GoToGrid - A Web-Oriented Tool in Support to Sites for LCG Installations	91
GraXML	92
Grand Challenges facing Storage Systems	92

Grid Collector: Using an Event Catalog to Speed up User Analysis in Distributed Environment	93
Grid Deployment Experiences: The path to a production quality LDAP based grid information system	93
Grid Enabled Analysis : Architecture, prototype and status	94
Grid Security	94
Grid2003 Monitoring, Metrics, and Grid Cataloging System	95
Grid3: An Application Grid Laboratory for Science	95
Griding The Nordic Supercomputing Infrastructure	96
Guidelines for Developing a Good GUI	96
H1OO - an analysis framework for H1	97
HEP Applications Experience with the European DataGrid Middleware and Testbed	97
HEP@HOME - A distributed computing system based on BOINC	98
HEPBook - A Personal Collaborative HEP notebook	99
Harp data and software migration from Objectivity to Oracle	100
Housing Metadata for the Common Physicist Using a Relational Database	100
How to build an event store - the new Kanga Event Store for BaBar	101
Huge Memory systems for data-intensive science	102
IGUANA Interactive Graphics Project: Recent Developments	102
INDICO - the software behind CHEP 2004	102
INTAS Discussion	103
IceTray: a Software Framework for IceCube	103
IgProf profiling tool	104
Implementation and Performance of the High-Level Trigger electron and photon selection for the ATLAS experiment at the LHC	104
Implementation of a reliable and expandable on-line storage for compute clusters	105
Improving Standard C++ for the Physics Community	106
InGRID - Installing GRID	106
InfiniBand for High Energy Physics	107
Information and Monitoring Services within a Grid Environment	107
Installing and Operating a Grid Infrastructure at DESY	108

Integrating Mutiple PC Farms into an uniform computing System with Maui	109
Integration of ATLAS Software in the Combined Beam Test	109
Interactive Data Analysis on the Grid using Globus 3 and JAS3	110
Ion transport simulation using Geant4 hadronic physics	110
JASSimApp plugin for JAS3: Interactive Geant4 GUI	111
JIM Deployment for the CDF Experiment	112
Jefferson Lab Data Acquisition Run Control System	112
Job Interactivity using a Steering Service in an Interactive Grid Analysis Environment	113
Job Monitoring in Interactive Grid Analysis Environment	113
Job Submission & Monitoring on the PHENIX Grid*	114
Job-monitoring over the Grid with GridIce infrastructure.	115
K5 @ INFN.IT: an infrastructure for the INFN cross REALM & AFS cell authentication.	115
LCG Conditions Database Project Overview	115
LCG Generator	116
LCIO persistency and data model for LC simulation and reconstruction	117
LEXOR, the LCG-2 Executor for the ATLAS DC2 Production System	117
LHC data files meet mass storage and networks: going after the lost performance	118
LambdaStation: A forwarding and admission control service to interface production network facilities with advanced research network paths	119
Lattice QCD Clusters at Fermilab	119
Lattice QCD Data and Metadata Archives at Fermilab and the International Lattice Data Grid	120
Linux for the CLEO-c Online system	120
MONARC2: A Processes Oriented, Discrete Event Simulation Framework for Modelling and Design of Large Scale Distributed Systems.	121
Managed Data Storage and Data Access Services for Data Grids	122
Managing software licences for a large research laboratory	122
Managing third-party software for the LCG	123
Mantis: the Geant4-based simulation specialization of the CMS COBRA framework	124
Mass Storage Management and the Grid	124
Methodologies and techniques for analysis of network flow data	125

Middleware for the next generation Grid infrastructure	125
Migrating PHENIX databases from object to relational model	126
Mis-use Cases for the Grid	126
MonALISA: An Agent Based, Dynamic Service System to Monitor, Control and Optimize Grid based Applications.	127
Monitoring CMS Tracker construction and data quality using a grid/web service based on a visualization tool	128
Monitoring a Petabyte Scale Storage System	128
Monitoring the CDF distributed computing farms	129
Monte Carlo Event Generation in a Multilanguage, Multiplatform Environment	129
Multi-Terabyte EIDE Disk Arrays running Linux RAID5	130
Muon Event Filter Software for the ATLAS Experiment at LHC	130
Muon Reconstruction Software in CMS	131
Network Architecture: lessons from the past, vision for the future	131
Network Information and Management Infrastructure Project	132
Networking for High Energy and Nuclear Physics as Global E-Science	133
New Applications of PAX in Physics Analyses at Hadron Colliders	133
New compact hierarchical mass storage system at Belle realizing a peta-scale system with inexpensive ice-raid disks and an S-ait tape library	134
New distributed offline processing scheme at Belle	134
New experiences with the ALICE High Level Trigger Data Transport Framework	135
New specific solids definitions in the Geant4 geometry modeller	136
Offline Software for the ATLAS Combined Test Beam	136
On Distributed Database Deployment for the LHC Experiments	137
On the Management of Certification Authority in Large Scale GRID Infrastructure	137
On-demand Layer VPN Support for Grid Applications	138
Online Monitoring and online calibration/reconstruction for the PHENIX experiment	139
OpenPAW	139
OpenScientist. Status of the project.	140
Operating the LCG and EGEE Production Grids for HEP	140

Operation of the CERN Managed Storage environment; current status and future directions	140
Optimizing Selection Performance on Scientific Data by utilizing Bitmap Indices	141
OptorSim: a Simulation Tool for Scheduling and Replica Optimisation in Data Grids	142
Overview and new developments in Geant4 electromagnetic physics	142
POOL Development Status and Plans	143
POOL Integration into three Experiment Software Frameworks	143
Panoramix	144
Parallel compilation of CMS software	144
Parallel implementation of Parton String Model event generator	145
Participation of Russian sites in the Data Challenge of ALICE experiment in 2004	145
Patching PCs	146
Paths: Specifying Multiple Job Outputs via Filter Expressions	146
Patriot: Physics Archives and Tools required to Investigate Our Theories	147
Performance analysis of Cluster File System on Linux	147
Performance of an operating High Energy Physics Data grid, D0SAR-grid	148
Performance of the ATLAS DAQ DataFlow system	149
Performance of the NorduGrid ARC and the Dulcinea Executor in ATLAS Data Challenge 2	150
Persistence for Analysis Objects	151
Physics Validation of the LHC Software	152
Physics validation of the simulation packages in a LHC-wide effort	152
Pixel Reconstruction in the CMS High-Level Trigger	152
Portable Gathering System for Monitoring and Online Calibration at Atlas	153
Porting CLEO software to Linux	153
Porting LCG Applications	154
Practical approaches to Grid workload and resource management in the EGEE project	154
Precision electromagnetic physics in Geant4: the atomic relaxation models	155
Precision validation of Geant4 electromagnetic physics	156
Predicting Resource Requirements of a Job Submission	156

Production Experience of the Storage Resource Broker in the BaBar Experiment	157
Production Management Software for the CMS Data Challenge	157
Production data export and archiving system for new data format of the BaBar experiment.	158
Production mode Data-Replication framework in STAR using the HRM Grid	158
Production of simulated events for the BaBar experiment by using LCG	159
PyBus – A Python Software Bus	160
Python-based physics analysis environment for LHCb	160
Quality Assurance and Testing in LCG	161
RDBC: ROOT DataBase Connectivity	161
ROOT : detector visualization	162
ROOT Graphical User Interface	162
Raw Ethernet based hybrid control system for the automatic control of suspended masses in gravitational waves interferometric detectors	163
Realization of a stable network flow with high performance communication in high bandwidth- delay product network	163
Recent Developments in the ROOT I/O	164
Recent evolutions of CMT. Multi-project and activity management.	164
Reflection-Based Python-C++ Bindings	165
Remote Shifting at the CLEO Experiment	166
Resource Predictors in HEP Applications	166
Results of the LHCb experiment Data Challenge 2004	167
Rio	168
Role of Tier-0, Tier-1 and Tier-2 Regional Centres in CMS DC04	168
Run II computing	169
SAMGrid Experiences with the Condor Technology in Run II Computing	170
SAMGrid Integration of SRMs	171
SAMGrid Monitoring Service and its Integration with MonALisa	171
SPHINX: A Scheduling Middleware for Data Intensive Applications on a Grid	172
SRB system at Belle/KEK	173
SRM AND GFAL TESTING FOR LCG2	173

ScotGrid: A prototype Tier 2 centre	174
Secure Grid Data Management Technologies in ATLAS	174
Self-Filling Histograms: An object-oriented analysis framework	175
Simplified deployment of an EDG/LCG cluster via LCFG-UML	175
Simulation and reconstruction of heavy ion collisions in the ATLAS detector.	176
Simulations and Prototyping of the LHCb L1 and HLT Triggers	176
Software Management in the HARP experiment	177
Software agents in data and workflow management	178
Software management infrastructure in the LCG Application Area	178
Specifying Selection Criteria using C++ Expression Templates	179
Status and Plans of the LCG PI Project	179
Status of the alignment calibrations in the ATLAS-Muon experiment	180
StoRM: grid middleware for disk resource management	180
Storage Resource Manager	181
Storage Resource Managers at Brookhaven	182
Summary - Computer Fabrics	182
Summary - Core Software	182
Summary - Distributed Computing Services	182
Summary - Distributed Computing Systems and Experiences	182
Summary - Event Processing	183
Summary - Online Computing	183
Summary - Wide Area Networking	183
Super scaling PROOF to very large clusters	183
Supporting the Development Process of the DataGrid Workload Management System Software with GNU autotools, CVS and RPM	184
Synergia: A Modern Tool for Accelerator Physics Simulation	184
Test of data transfer over an international network with a large RTT	185
Test of the ATLAS Inner Detector reconstruction software using combined test beam data	186
Testbed Management for the ATLAS TDAQ	186
Testing the CDF Distributed Computing Framework	187

The ALICE Data Challenge 2004 and the ALICE distributed analysis prototype	187
The ALICE Experiment Control System	188
The ALICE High Level Trigger	188
The ARDA Prototypes	189
The ATLAS Computing Model	190
The ATLAS DAQ system	190
The AliRoot framework, status and perspectives	191
The Architecture of the ZEUS Second Level Global Tracking Trigger	191
The Athena Control Framework in Production, New Developments and Lessons Learned	192
The Atlantis event visualisation program for the ATLAS experiment	192
The BABAR Analysis Task Manager	193
The BIRN Project: Distributed Information Infrastructure and Multi-scale Imaging of the Nervous System (BIRN = Biomedical Informatics Research Network)	194
The Binary Cascade	194
The CEDAR Project	195
The CMS User Analysis Farm at Fermilab	196
The Clarens Grid-enabled Web Services Framework: Services and Implementation	196
The Condor based CDF CAF	197
The Configurations Database Challenge in the ATLAS DAQ System	197
The D0 Virtual Center and planning for large scale computing	198
The DAQ system for the Fluorescence Detectors of the Pierre Auger Observatory	199
The DZERO Run II Level 3 Trigger and Data Acquisition System	200
The Description of the Atlas Detector	200
The Design, Installation and Management of a Tera-Scale High Throughput Cluster for Particle Physics Research	201
The Evolution of Computing: Slowing down? Not Yet!	201
The Evolving Wide Area Network Infrastructure in the LHC era	202
The FEDRA - Framework for Emulsion Data Reconstruction and Analysis in OPERA exper- iment.	202
The FreeHEP Java Library Root IO package	203
The GANGA user interface interface for physics analysis on distributed resources	203

The GeoModel Toolkit for Detector Description	204
The Geometry Package for the Pierre Auger Observatory	204
The GridSite authorization system	205
The High Level Filter of the H1 Experiment at HERA	205
The High Level Trigger software for the CMS experiment	206
The IBM Research Global Technology Outlook	206
The LCG Project - Preparing for Startup	207
The LCG Savannah software development portal	207
The LCG-AliEn interface, a realization of a MetaGrid system	207
The LHCb Configuration Database	208
The Network Security Protection System at IHEP-Net	208
The Next Generation Root File Server	209
The Offline Framework of the Pierre Auger Observatory	210
The Open Science Grid (OSG)	210
The PHENIX Event Builder	211
The Project CampusGrid	212
The ROOT Linear Algebra	212
The SAMGrid Database Server Component: Its Upgraded Infrastructure and Future Development Path	212
The SEAL C++ Reflection System	213
The SEAL Component Model	214
The STAR Unifid Meta-Scheduler project, a front end around evolving technologies for user analysis and data production.	214
The TAG Collector - A Tool for Atlas Code Release Management	215
The Track Extrapolation package in the new ATLAS Tracking Realm	215
The Virtual Geometry Model	216
The Virtual MonteCarlo : status and applications	216
The Virtual Organization Membership Service eXtension project (VOX)	217
The WEB interface for the ATLAS/LCG MySQL Conditions Databases and performance constraints in the visualisation of extensive scientific/technical data	218
The ZEUS Global Tracking Trigger	219

The ZEUS Global Tracking Trigger Barrel Algorithm	219
The ZOOM Minimization Package	219
The application of PowerPC/VxWorks to the read-out subsystem of the BESIII DAQ	220
The architecture of the AliEn system	221
The deployment mechanisms for the ATLAS software.	221
The evolution of the distributed Event Reconstruction Control System in BaBar	222
The impact of e-science	222
The introduction to BES computing environment	223
The new BaBar Analysis Model	223
The role of legacy services within ATLAS DC2	224
The role of scientific middleware in the future of HEP computing	224
The simulation for the ATLAS experiment: present status and outlook	225
The status of Fermilab Enstore Data Storage System	225
Tier-1 and Tier-2 Real-time Analysis experience in CMS Data Challenge 04	226
Tools for GRID deployment of CDF offline and SAM data handling systems for Summer 2004 computing.	227
Toward a Grid Technology Independent Programming Interface for HEP Applications	228
Track reconstruction in high density environment	228
Tracking of long lived hyperons in silicon detector at CDF.	229
Transparently managing time varying conditions and detector data on ATLAS.	229
Update On the Status of the FLUKA Monte Carlo Transport Code	230
Usage of ALICE Grid middleware for medical applications	231
Usage statistics and usage patterns on the NorduGrid: Analyzing the logging information collected on one of the largest production Grids of the world.	232
Use of Condor and GLOW for CMS Simulation Production	232
Using HEP Systems to Provide Storage for Biologists	233
Using Nagios for intrusion detection	234
Using Tripwire to check cluster system integrity	234
Using the reconstruction software, ORCA, in the CMS datachallenge	235
Validation of the GEANT4 Bertini Cascade model and data analysis using the Parallel ROOT Facility on a Linux cluster	235

Vertex finding and B-tagging algorithms for the ATLAS Inner Detector	236
Visualisation of the ROOT geometries (TVolume and TGeoVolume) with Coin-based model.	236
Volume-based representation o the magnetic field	237
WAN Emulation Development and Testing at Fermilab	238
WIRED 4 - A generic Event Display plugin for JAS 3	238
Welcome Address	239
Wide Area Network Monitoring system for HEP experiments at Fermilab	239
Writing Extension Modules (Plug-ins) for JAS3	240
XML I/O in ROOT	240
XNetFile, a fault tolerant extension of ROOT TNetFile	241
dCache, LCG Storage Element and enhanced use cases	241

Poster Session 3 - Tableau: 11 / 269

”RecPack”, a general reconstruction tool-kit

A. CERVERA VILLANUEVA¹ ; J. A. HERNANDO MORATA² ; J. J. GOMEZ CADENAS³

¹ *University of Geneva*

² *CERN and University of Santiago de Compostela*

³ *University of Valencia*

Auteur(s) contact: anselmo.cervera.villanueva@cern.ch

We have developed a c++ software package, called “RecPack”, which allows the reconstruction of dynamic trajectories in any experimental setup. The basic utility of the package is the fitting of trajectories in the presence of random and systematic perturbations to the system (multiple scattering, energy loss, inhomogeneous magnetic fields, etc) via a Kalman Filter fit. It also includes an analytical navigator which allows: extrapolation of the trajectory parameters (and their covariance matrix) to any surface, path length computations, matching functions (trajectory-trajectory, trajectory-measurement, etc) and much more. The RecPack tool-kit also includes the algorithms for vertex fitting via Kalman Filter, and the necessary tools for easily coding pattern recognition algorithms. In summary, “RecPack” provides all the necessary tools and algorithms that are common to any reconstruction program. In addition, a toy simulator is provided. This is very useful to debug new reconstruction algorithms and also to perform simple physics analysis. The modularity of the package allows extensions in any direction: new propagation models, measurement types, volume and surface types, fitting algorithms, etc.

I

Plenary / 521

50 years of Computing at CERN

D. WILLIAMS^{None}

“Where are your Wares”

Computing in the broadest sense has a long history, and Babbage (1791-1871), Hollerith (1860-1929) Zuse (1910-1995), many other early pioneers, and the wartime code breakers, all made important breakthroughs. CERN was founded as the first valve-based digital computers were coming onto the market.

I will consider 50 years of Computing at CERN from the following viewpoints:-

Where did we come from? What happened? Who was involved? Which wares (hardware, software, netware, peopleware and now middleware) were important? Where did computers (not) end up in a physics lab? What has been the impact of computing on particle physics? What about the impact of particle physics computing on other sciences? And the impact of our computing outside the scientific realm?

I hope to conclude by looking at where we are going, and by reflecting on why computing is likely to remain challenging for a long time yet.

The topic is so vast that my remarks are likely to be either prejudiced or trivial, or both.

Computer Fabrics / 138**64-Bit Opteron systems in High Energy and Astroparticle Physics**

P. Wegner¹ ; S. Wiesand¹

¹ DESY

Auteur(s) contact: stephan.wiesand@desy.de, peter.wegner@desy.de

64-Bit commodity clusters and farms based on AMD technology meanwhile have been proven to achieve a high computing power in many scientific applications. This report first gives a short introduction into the specialties of the amd64 architecture and the characteristics of two-way Opteron systems.

Then results from measuring the performance and the behavior of such systems in various Particle Physics applications as compared to the classical 32-Bit systems are presented. The investigations cover analysis tools like ROOT, Astrophysics simulations based on CORSIKA and event reconstruction programs. Another field of investigations are parallel high performance clusters for Lattice QCD calculations, and n-loop calculations based on perturbative methods in quantum field theory using the formula manipulation program FORM.

In addition to the performance results the compatibility of 32- and 64-Bit architectures and Linux operating system issues, as well as the impact on fabric management are discussed.

It is shown that for most of the considered applications the recently available 64-bit commodity computers from AMD are a viable alternative to comparable 32-Bit systems.

Online Computing / 254**A Control Software for the ALICE High Level Trigger**

H. Tilsner¹ ; T.M. Steinbeck¹ ; V. Lindenstruth¹

¹ KIRCHHOFF INSTITUTE OF PHYSICS, RUPRECHT-KARLS-UNIVERSITY HEIDELBERG, for the Alice Collaboration

Auteur(s) contact: timm.steinbeck@kip.uni-heidelberg.de

The Alice High Level Trigger (HLT) cluster is foreseen to consist of 400 to 500 dual SMP PCs at the start-up of the experiment. The software running on these PCs will consist of components communicating via a defined interface, allowing flexible software configurations. During Alice's operation the HLT has to be continuously active to avoid detector dead time. To ensure that the several hundred software components, distributed throughout the cluster, operate and interact properly, a control software was written that is presented here. It was designed to run distributed over the cluster and to support control program hierarchies. Distributed operation avoids central performance bottlenecks and single-points-of-failures. The last point is of particular importance, as each of the commodity type PCs in the HLT cluster cannot be relied upon to operate continuously. Control hierarchies in turn are relevant for scalability over the required number of nodes. The software makes use of existing and widely used technologies: Configurations of programs to be controlled are saved in XML, while Python is used as a scripting language and to specify actions to execute. Interface libraries are used to access the controlled components, presenting a uniform interface to the control program. Using these mechanisms the control software remains generic and can

be used for other purposes as well. It is being used for HLT data challenges in Heidelberg and is planned for use during upcoming beam tests.

Distributed Computing Systems and Experiences / 486

A Data Grid for the Analysis of Data from the Belle Experiment

G R. Moloney^{None} ; J. Smillie^{None} ; L J. Winton^{None} ; S. McMahon^{None}

Auteur(s) contact: glenn@physics.unimelb.edu.au

We have developed and deployed a data grid for the processing of data from the Belle experiment, and for the production of simulated Belle data. The Belle Analysis Data Grid brings together compute and storage resources across five separate partners in Australia, and the Computing Research Centre at the KEK laboratory in Tsukuba, Japan.

The data processing resources are general purpose, shared use, compute clusters at the Universities of Melbourne and Sydney, the Australian Partnership for Advanced Computing (APAC), the Victorian Partnership for Advanced Computing (VPAC) and the Australian Centre for Advanced Computing and Communications (AC3).

This system is in use for the Australian contribution to the production of simulated data for the Belle experiment, and for physics analyses.

The Storage Resource Broker (SRB), from the San Diego Supercomputing Centre, is used to provide a robust underlying data repository. A federation of SRB servers has been established to share and manage Belle data between the KEK laboratory, the mass data store at the Australian National University (ANU) and satellite storage at each of the compute clusters.

The globus toolkit is the underlying technology for the management of the computing resources, and the despatching of jobs. A network aware job scheduler has been developed. The scheduler queries the SRB servers for location of data replicas, and arranges scheduling of processing and production jobs on the compute resources according to a static model of the network connectivity and dynamic assessment of the relative system loads.

Core Software / 340

A Dynamically Reconfigurable Data Stream Processing System

J. Nogiec¹ ; K Trombly-Freytag²

¹ *FERMI NATIONAL ACCELERATOR LABORATORY*

² *Fermilab*

Auteur(s) contact: nogiec@fnal.gov

The paper describes a component-based framework for data stream processing that allows for configuration, tailoring, and run-time system reconfiguration. The

system's architecture is based on a pipes and filters pattern, where data is passed through routes between components. Components process data and add, substitute, and/or remove named data items from a data stream. They can also manipulate data streams by buffering data, compressing/decompressing individual streams, and combining, splitting, or synchronizing multiple data streams. Configurable general-purpose filters for manipulating streams, visualizing data, persisting data, and reading data from various standard data sources are supplemented with many application specific filters, such as DSP, scripting, or instrumentation-specific components. A network of pipes and filters can be dynamically reconfigured at run-time, in response to a preplanned sequence of processing steps, operator intervention, or a change in one or more data streams. Four distinctive methods supporting reconfiguration are provided by the framework: modification of data routes, management of components' activity states, triggering processing based on the content of the data, or the use of source addressing in components. The framework can be used to build static data stream processing applications such as monitoring or data acquisition systems as well as self-adjusting systems that would adapt their processing algorithm, presentation layer, or data persistency layer in response to changes in input data streams.

Distributed Computing Systems and Experiences / 309

A GRID approach for Gravitational Waves Signal Analysis with a Multi-Standard Farm Prototype

A. Eleuteri¹ ; F. Acernese¹ ; F. Barone² ; F.M. Taurino³ ; G. Russo⁴ ; G. Tortone³ ; L. Giordano⁵ ; L. Milano⁴ ; P. Mastroserio⁶ ; R. De Rosa⁴ ; R. Esposito⁶ ; S. Pardi⁷

¹ *Dipartimento di Scienze Fisiche - Università di Napoli Federico II*

² *Dipartimento di Scienze Farmaceutiche - Università di Salerno*

³ *INFN, Napoli*

⁴ *Dipartimento di Scienze Fisiche - Università di Napoli Federico II*

⁵ *INFN and University of Napoli*

⁶ *INFN - sez.Napoli*

⁷ *DIPARTIMENTO DI MATEMATICA ED APPLICAZIONI "R.CACCIOPPOLI"*

Auteur(s) contact: spardi@na.infn.it

The standard procedures for the extraction of gravitational wave signals coming from coalescing binaries provided by the output signal of an interferometric antenna may require computing powers generally not available in a single computing centre or laboratory. A way to overcome this problem consists in using the computing power available in different places as a single geographically distributed computing system. This solution is now effective within the GRID environment, that allows distributing the required computing effort for specific data analysis procedure among different sites according to the available computing power.

Within this environment we developed a system prototype with application software for the experimental tests of a geographically distributed computing system for the analysis of gravitational wave signal from coalescing binary systems. The facility has been developed as a general purpose system that uses only standard hardware and software components, so that it can be easily upgraded and configured. In fact, it can be partially or totally configured as a GRID farm, as MOSIX farm or as MPI farm. All these three configurations may coexist since the facility can be split into configuration subsets. A full description of this farm is reported, together with the results of the performance tests and planned developments.

Event Processing / 100**A Gaussian-sum filter for vertex reconstruction**R. Fruehwirth¹ ; T. Speer²¹ *Institut für Hochenergiephysik der OAW, Vienna, Austria*² *UNIVERSITY OF ZURICH, SWITZERLAND***Auteur(s) contact:** thomas.speer@cern.ch

A vertex fit algorithm was developed based on the Gaussian-sum filter (GSF) and implemented in the framework of the CMS reconstruction program. While linear least-squares estimators are optimal in case all observation errors are Gaussian distributed, the GSF offers a better treatment of the non-Gaussian distribution of track parameter errors when these are modeled by Gaussian mixtures.

In addition, when using electron tracks reconstructed with an electron-reconstruction Gaussian-sum filter, the full mixture can be used rather than the approximation by a single Gaussian.

Properties, results and performance of this filter with simulated data will be shown, and compared to the Kalman filter and to robust filters.

Wide Area Networking / 88**A Globally Distributed Real Time Infrastructure for World Wide Collaborations**D. Adamczyk¹ ; D. Lattka¹ ; G. Denis¹ ; H. Newman¹ ; I. Legrand² ; J. Fernandes³ ; J. SUCIK¹ ; K. Wei¹ ; P. FARKAS¹ ; P. Galvez²¹ *Caltech*² *CALTECH*³ *Caltech/CERN***Auteur(s) contact:** galvez@caltech.edu

VRVS (Virtual Room Videoconferencing System) is a unique, globally scalable next-generation system for real-time collaboration by small workgroups, medium and large teams engaged in research, education and outreach. VRVS operates over an ensemble of national and international networks. Since it went into production service in early 1997, VRVS has become a standard part of the toolset used daily by a large sector of HENP, and it is used increasingly for other DoE/NSF-supported programs. Today, the VRVS Web-based system is regularly accessed by more than 30,000 registered hosts running the VRVS software in more than 103 countries. There are currently 78 VRVS “reflectors” that create the interconnections and manage the traffic flow, in the Americas, Europe and Asia. New reflectors recently have been installed in Brazil, China, Pakistan, Australia and Slovakia.

VRVS is global in scope: it covers the full range of existing and emerging protocols and the full range of client devices for collaboration, from mobile systems through desktops to installations in large auditoria. VRVS will be integrated with the Grid-enabled Analysis Environment (GAE) now under development at Caltech in partnership with the GriPhyN, iVDGL and PPDG projects in the US, and Grid projects in Europe.

A major architectural change is currently in development. The new version v4.0, is expected to be deployed in early 2005. We will describe the current operational state of the VRVS service and provide a description of the new architecture including all the new and advanced functionalities that will be added.

Online Computing / 223

A Hardware Based Cluster Control And Management System

H. TILSNER¹ ; L. Hess² ; R. Panse² ; T. Alt² ; V. Lindenstruth²

¹ *Kirchhoff Institute of Physics, Heidelberg, Germany*

² *KIRCHHOFF INSTITUTE FOR PHYSICS - UNIVERSITY OF HEIDELBERG*

Auteur(s) contact: panse@kip.uni-heidelberg.de

Super-computers will be replaced more and more by PC cluster systems. Also future LHC experiments will use large PC clusters. These clusters will consist of off-the-shelf PCs, which in general are not built to run in a PC farm. Configuring, monitoring and controlling such clusters requires a serious amount of time consuming and administrative effort.

We propose a cheap and easy hardware solution for this issue. The main item of our cluster control system is the Cluster Interface Agent card (CIA).

The CIA card is a low-cost PCI expansion card equipped with a network interface. With the aid of the CIA card the computer can be fully controlled remotely, independent of the state of the node itself. The card combines a number of feature needed for this remote control, including power management and reset. The card operates entirely independent of the PC and can remain powered while the PC may even be powered down. It offers a wide range of automatization features, including automatic installation of the operating system, changing BIOS settings or booting a rescue disk and also to monitor and debug the node. With the aid of PCI scans and hardware tests errors and pending failures can be easily detected in an early stage.

Working prototypes exist. The presentation will outline the status of the project and first implementation results of the preproduction devices, currently being built.

Online Computing / 482

A Level-2 trigger algorithm for the identification of muons in the Atlas Muon Spectrometer

A. Di Mattia¹

¹ *INFN*

The Atlas Level-2 trigger provides a software-based event selection after the initial Level-1 hardware trigger. For the muon events, the selection is decomposed in a number of broad steps: first, the Muon Spectrometer data are processed to give physics quantities associated to the muon track (standalone features extraction) then,

other detector data are used to refine the extracted features. The “muFast” algorithm performs the standalone feature extraction, providing a first reduction of the muon event rate from Level-1. It confirms muon track candidates with a precise measurement of the muon momentum. The algorithm is designed to be both conceptually simple and fast so as to be readily implemented in the demanding online environment in which the Level-2 selection code will run. Never-the-less its physics performance approaches, in some cases, those of the offline reconstruction algorithms. This paper describes the implemented algorithm together with the software techniques employed to increase its timing performance.

Distributed Computing Services / 388

A Lightweight Monitoring and Accounting System for LHCb DC’04 Production

J. Saborido-Silva¹ ; M. Sanchez-Garcia² ; R. Graciani-Diaz³ ; R. Vizcaya Carrillo³ ; V. Garonne⁴

¹ *University of Santiago de Compostela*

² *UNIVERSITY OF SANTIAGO DE COMPOSTELA*

³ *University of Barcelona*

⁴ *IN2P3*

Auteur(s) contact: manusan@usc.es

The LHCb Data Challenge 04 includes the simulation of over 200 M simulated events using distributed computing resources on N sites and extending along 3 months. To achieve this goal a dedicated Production grid (DIRAC) has been deployed. We will present the Job Monitoring and Accounting services developed to follow the status of the production along its way and to evaluate the results at the end of the Data Challenge.

The end user connects with a web browser to WEB-SERVER applications showing dynamic reports for a whole set of possible queries. These applications in turn interrogate the Job Monitoring Service of the DIRAC Workload Management system and Accounting Database service by means of dedicated XML-RPC interfaces, querying for the information requested by the user. The reports provide an uniform view of the usage of the computing resources available. All the system components are implemented as a set of cooperating python classes following the design choice of LHCb. The different services are distributed over a number of independent machines. This allows to achieve the scalability level of multiple thousands of concurrent jobs monitored by the system

Poster Session 2 - Tableau: 51 / 296

A Multidimensional Approach to the Analysis of Grid Monitoring Data

G. Rubini¹ ; G. Tortone² ; N. De Bortoli² ; S. Andreozzi¹ ; S. Fantinel³

¹ *INFN-CNAF*

² *INFN-Napoli*

³ *INFN-LNL*

Auteur(s) contact: gianluca.rubini@cnafe.infn.it

Analyzing Grid monitoring data requires the capability of dealing with multidimensional concepts intrinsic to Grid systems. The meaningful dimensions identified in recent works are the physical dimension referring to geographical location of resources, the Virtual Organization (VO) dimension, the time dimension and the monitoring metrics dimension. In this paper, we discuss the application of On-Line Analytical Processing (OLAP), an approach to the fast analysis of shared multidimensional information, to the mentioned problem. OLAP relies on structures called 'OLAP cubes', that are created by a reorganization of data contained inside a relational database, thus transforming operational data into dimensional data.

Our OLAP model is a four-dimension cuboid based on time, geographic, Virtual Organization (VO), and monitoring metric. Time and geographic dimensions have total order relation and form two concept hierarchies, respectively hours<days<weeks<months and services<hosts<sites<countries. This model is applied to a set of Grid monitoring data generated by a monitoring tool called GridICE.

Event Processing / 469

A New STAR Event Reconstruction Chain

A. Rose¹ ; B. Hippolyte² ; C. Pruneau¹ ; J. Lauret³ ; M. Calderon⁴

¹ *WAYNE STATE UNIVERSITY*

² *Yale University*

³ *Brookhaven National Laboratory*

⁴ *Brookhaven National Lab*

Auteur(s) contact: pruneau@physics.wayne.edu

We present the design and performance analysis of a new event reconstruction chain deployed for analysis of STAR data acquired during the 2004 run and beyond. The creation of this new chain involved the elimination of obsolete FORTRAN components, and the development of equivalent or superior modules written in C++.

The new reconstruction chain features a new and fast TPC cluster finder, a new track reconstruction software (ITTF discussed at CHEP2003), which seamlessly integrate all detector components of the experiment, a new vertex finder, and various post-tracking analysis modules including a V0 finder, and a track kink finder. The new chain is the culmination of a large software development effort involving in excess of ten FTEs.

Poster Session 3 - Tableau: 15 / 270

A Pattern-based Continuous Integration Framework for Distributed EGEE Grid Middleware Development

Auteur(s): A. Di Meglio¹

Co-auteur(s) E. Ronchieri² ; J. Flammer¹ ; M. Zurek¹ ; R. Harakaly¹

¹ CERN

² INFN

Auteur(s) contact: alberto.di.meglio@cern.ch

Software Configuration Management (SCM) Patterns and the Continuous Integration method are recent and powerful techniques to enforce a common software engineering process across large, heterogeneous, rapidly changing development projects where a rapid release lifecycle is required. In particular the Continuous Integration method allows tracking and addressing problems in the software components integration as early as possible in the release cycle. Since new incremental code builds are done several times per day, only small amounts of new code is built and integrated at relatively short intervals. Developers are immediately notified of arising problems and integrators can pinpoint configuration and build problems to the level of single files within any given software component. This paper presents the implementation and the initial results of the application of such techniques in the SCM and Integration of the EGEE Grid Middleware software. The software is based on a Service Oriented Architecture model where services are developed in different programming languages by development groups in several European locations under stringent quality requirements. A number of basic SCM patterns, such as the Workspace, the Active Line, the Repository, are introduced and the Continuous Integration tools used in the project are presented with a discussion of the advantages and disadvantages of using the method.

Computer Fabrics / 330

A Regional Analysis Center at the University of Florida

C. Prescott¹ ; D. Bourilkov¹ ; J. Rodriguez¹ ; P. Avery¹ ; R. Cavanaugh¹ ; Y. Fu¹

¹ UNIVERSITY OF FLORIDA

Auteur(s) contact: richard.cavanaugh@cern.ch

The High Energy Physics Group at the University of Florida is involved in a variety of projects ranging from High Energy Experiments at hadron and electron positron colliders to cutting edge computer science experiments focused on grid computing. In support of these activities members of the Florida group have developed and deployed a local computational facility which consists of several service nodes, computational clusters and disk storage services. The resources contribute collectively or individually to a variety of production and development activities such as the UFlorida Tier2 center for the CMS experiment at the Large Hadron Collider (LHC), Monte Carlo production for the CDF experiment at Fermi Lab, the CLEO experiment, and research on grid computing for the GriPhyN and iVDGL projects. The entire collection of servers, clusters and storage services is managed as a single facility using the ROCKS cluster management system. Managing the facility as a single centrally managed system enhances our ability to relocate and reconfigure the resources as necessary in support of both research and production activities. In this paper we describe the architecture deployed, including details on our local implementation of the ROCKS systems, how this simplifies the maintenance and administration of the facility and finally the advantages and disadvantages of using such a scheme to manage a modest size facility.

Grid Security / 122

A Scalable Grid User Management System for Large Virtual Organization

D. Yu¹ ; G. Carcassi¹ ; G. Smith¹ ; J. Smith¹ ; J. Spiletic¹ ; T. Carter¹ ; T. Wlodek¹ ; X. Zhao¹ ; Z. Liu¹¹ *BROOKHAVEN NATIONAL LABORATORY***Auteur(s) contact:** carcassi@bnl.gov, dtyu@bnl.gov

We present a work-in-progress system, called GUMS, which automates the processes of Grid user registration and management and supports policy-aware authorization at well. GUMS builds on existing VO management tools (LDAP VO, VOMS and VOMRS) with a local grid user management system and a site database which stores user credentials, accounting history and policies in XML format. We use VOMRS, being developed by Fermilab, to collect user information and register legitimate users into the VOMS server.

Our local grid user management system jointly retrieves user information and VO policies from multiple VO databases based on site security policies. Authorization can be done by mapping the user's credential to local accounts. Four different mapping schemes have been implemented: user's existing account, recyclable pool account, non-recyclable pool account and group shared account. The mapping selection is determined by the type of target resource and its usage policies. We already deployed our automatic grid mapfile generators on the BNL Grid Gatekeeper, GridFtp server and HPSS mass storage system. Work is in progress to enable "single-sign-on" based upon X509 certificate credential for job execution and access to both disk and tape storage resources.

Poster Session 3 - Tableau: 1 / 225

A database perspective on CMS detector data

A.T.M. Aerts¹ ; F. Glege² ; M. Liendl²¹ *TU Eindhoven*² *CERN***Auteur(s) contact:** martin.liendl@cern.ch, a.t.m.aerts@tue.nl

Building a state of the art high energy physics detector like CMS requires strict interoperability and coherency in the design and construction of all sub-systems comprising the detector. This issue is especially critical for the many database components that are planned for storage of the various categories of data related to the construction, operation, and maintenance of the detector like event data, slow control data, conditions data, calibration data, event meta data, etc The data structures needed to operate the detector as a whole need to be present in the database before the data is entered. Changing these structures for a database system that already contains a substantial amount of data is a very time and labour consuming exercise that needs to be avoided. Cases where the detector needs to be treated as a whole are detector operation (control, error tracking, conditions) and the interfacing of there construction and simulation software.

In this paper we propose to use the detector geometry as the structure connecting the various elements. The design and implementation of a relational database that captures the CMS detector geometry and the detector components is discussed. The detector geometry can serve as a core component in several other databases in order to make them interoperable. It also provides a common viewpoint between the physical detector and its image in the reconstruction software. Some of the necessary extensions to the detector description are discussed.

Poster Session 1 - Tableau: 33 / 180

A database prototype for managing computer systems configurations

N. Sinanis¹ ; Z. Toteva²

¹ *CERN/CMS*

² *Sofia University/CERN/CMS*

Auteur(s) contact: zhechka.toteva@cern.ch, nick.sinanis@cern.ch

We describe a database solution in a web application to centrally manage the configuration information of computer systems. It extends the modular cluster management tool Quattor with a user friendly web interface.

System configurations managed by Quattor are described with the aid of PAN, a declarative language with a command line and a compiler interface. Using a relational schema, we are able to build a database for efficient data storage and configuration data processing. The relational schema ensures the consistency of the described model while the standard database interface ensures the fast retrieval of configuration information and statistic data.

The web interface simplifies the typical administration and routine operations tasks, e.g. definition of new types, configuration comparisons and updates etc. We present a prototype built on the above ideas and used to manage a cluster of developer workstations and specialised services in CMS.

Distributed Computing Systems and Experiences / 43

A distributed, Grid-based analysis system for the MAGIC telescope

A. Forti¹ ; A. deAngelis¹ ; C. Bigongiari² ; G. Cabras¹ ; H. Kornmayer³ ; M. Delfino⁴ ; M. Frailis¹ ; M. Hardt³ ; M. Kunze³ ; M. Mazzucato⁵ ; M. Piraccini¹

¹ *University of Udine*

² *INFN, Padua*

³ *FORSCHUNGSZENTRUM KARLSRUHE (FZK)*

⁴ *PIC, Barcelona*

⁵ *INFN-Padova*

The observation of high-energetic gamma-rays with ground based air cerenkov telescopes is one of the most exciting areas in modern astro particle physics. End of the year 2003 the MAGIC telescope started operation. The low energy threshold for gamma-rays together with different background

sources leads to a considerable amount of data. The analysis will be done in different institutes spread over Europe. The production of Monte Carlo events including the simulation of Cerenkov light in the atmosphere is very computing intensive and another challenge for a collaboration like MAGIC. Therefore the MAGIC telescope collaborations will take the opportunity to use Grid technology to set up a distributed computational and data intensive analysis system with nowadays available technology. The basic architecture of such a distributed, Europe wide Grid system will be presented. First implementation results will be shown. This Grid might be the starting point for a wider distributed astro particle Grid in Europe.

Poster Session 2 - Tableau: 19 / 243

A general and flexible framework for virtual organization application tests in a grid system

A. Pierro¹ ; G. Donvito¹ ; G. Maggi¹ ; T. Coviello¹

¹ *INFN Via E. Orabona 4 I - 70126 Bari Italy*

Auteur(s) contact: tommaso.coviello@ba.infn.it

A grid system is a set of heterogeneous computational and storage resources, distributed on a large geographic scale, which belong to different administrative domains and serve several different scientific communities named Virtual Organizations (VOs). A virtual organization is a group of people or institutions which collaborate to achieve common objectives. Therefore such system has to guarantee the coexistence of different VO's applications providing them the suitable run-time environment. Hence tools are needed both at local and central level for testing and detecting eventually bad software configuration on a grid site.

In this paper we present a web based tool which permits to a Grid Operational Centre (GOC) or a Site Manager to test a grid site from the VO viewpoint.

The aim is to create a central repository for collecting both existing and emerging VO tests. Each VO test may include one or more specific application tests, and each test could include one or more subtests, arranged in a hierarchic structure.

A general and flexible framework is presented capable to include VO tests straightforwardly by means of a description file. Submission of a bunch of tests to a particular grid site is made available through a web portal. On the same portal, past and current results and logs can be browsed.

Poster Session 3 - Tableau: 19 / 483

A high-level language for specifying detector coincidences

F. Gray¹

¹ *UNIVERSITY OF CALIFORNIA, BERKELEY*

Auteur(s) contact: fgray@berkeley.edu

The muCap experiment at the Paul Scherrer Institut (PSI) will measure the rate of muon capture on the proton to a precision of 1% by comparing the apparent lifetimes of positive and negative muons in hydrogen. This rate may be related to the induced pseudoscalar weak form factor of the proton.

Superficially, the muCap apparatus looks something like a miniature model of a collider detector. Muons pass through several beam counters before reaching a hydrogen-filled time projection chamber (TPC) at its core, which acts as both a stopping target and the primary muon detector. It is surrounded by cylindrical wire chambers and a scintillator hodoscope to observe the Michel electrons that emerge from muon decay. The first key step in the analysis of our data is the proper definition of coincidence events across these many detector layers, maximizing the signal significance by suppressing accidental and pileup backgrounds. Part of our analysis software is written in a special-purpose high-level language, called “muon query language” (MQL), in which these coincidences may be specified cleanly. It uses a variant of the relational model, representing the data as a set of tables upon which selection and join operations may be performed. ROOT histograms and trees are defined based on the contents of tables. A preprocessor generates optimized C++ code that implements the operations described in the MQL file, which is suitable for incorporation into our analyzer framework. This talk will describe the MQL approach and our collaboration’s experience with it.

Event Processing / 75

A kinematic and a decay chain reconstruction library

K. Prokofiev¹ ; T. Speer¹

¹ *Physik-Institut Universitat Zurich*

Auteur(s) contact: kirill.prokofiev@cern.ch

A kinematic fit package was developed based on Least Means Squared minimization with Lagrange multipliers and Kalman filter techniques and implemented in the framework of the CMS reconstruction program. The package allows full decay chain reconstruction from final state to primary vertex according to the given decay model. The class framework allowing decay tree description on every reconstruction step will be described in details. Package extension to any type of physics object reconstructed in CMS, integration to general CMS reconstruction framework and related questions will be discussed. Examples of decay chain models, constraints and their application on Bs reconstruction will be presented.

Event Processing / 333

A statistical toolkit for data analysis

A. Pfeiffer¹ ; A. Ribon¹ ; B. Mascialino² ; M.G. Pia³ ; P. VIARENGO⁴ ; S. Donadio² ; S. Guatelli⁵

¹ *CERN*

² *INFN Genova*

³ *INFN GENOVA*

⁴ *IST, Genova, Italy*

⁵ *INFN Genova, Italy*

Auteur(s) contact: maria.grazia.pia@cern.ch

Statistical methods play a significant role throughout the life-cycle of HEP experiments, being an essential component of physics analysis. We present a project in progress for the development of an object-oriented software toolkit for statistical data analysis. More in particular, the Statistical Comparison component of the toolkit provides algorithms for the comparison of data distributions in a variety of use cases typical of HEP experiments, as regression testing (in various phases of the software life-cycle), validation of simulation through comparison to experimental data, comparison of expected versus reconstructed distributions, comparison of data from different sources - such as different sets of experimental data, or experimental with respect to theoretical distributions. The toolkit contains a variety of goodness-of-fit tests, from chi-squared to Kolmogorov-Smirnov, to less known, but generally much more powerful tests such as Anderson-Darling, Cramer-von Mises, Kuiper, Tiku etc.

Thanks to the component-based design and the usage of the standard AIDA interfaces, this tool can be used by other data analysis systems or integrated in experimental software frameworks. We present the architecture of the system, the statistics methods implemented and some results of its applications to the comparison of Geant4 simulations with respect to experiment.

Core Software / 397

AIDA, JAIDA and AIDAJNI: Data Analysis using interfaces

M. Donszelmann¹ ; M. Turri¹ ; T. Johnson¹ ; V. Serbo¹

¹ SLAC

Auteur(s) contact: mark.donszelmann@slac.stanford.edu, serbo@slac.stanford.edu

AIDA, Abstract Interfaces for Data Analysis, is a set of abstract interfaces for data analysis components: Histograms, Ntuples, Functions, Fitter, Plotter and other typical analysis categories. The interfaces are currently defined in Java, C++ and Python and implementations exist in the form of libraries and tools using C++ (Anaphe/Lizard, OpenScientist), Java (Java Analysis Studio) and Python (PAIDA).

JAIDA is the full implementation of AIDA in Java. It is used internally by JAS3 as its analysis core but it can also be used independently for either batch or interactive processing, or for web applications to access data, make plots and simple data analysis through a browser. Some of the JAIDA features are the ability to open AIDA, ROOT and PAW files and the support of an extensible set of fit methods (chi-square, least squares, binned/unbinned likelihood, etc) to be matched with an extensible set of optimizers including Minuit and Uncmin.

AIDAJNI is glue code between C++ and Java that allows any C++ code to access any Java implementation of the AIDA interfaces. For example AIDAJNI is used with Geant4 to access the JAIDA implementation of AIDA.

This paper gives an update on the AIDA 3.2.1 interfaces and its corresponding JAIDA implementation. Examples will be provided on how to use JAIDA within JAS3, as a standalone library and from C++ using AIDAJNI.

References:

<http://aida.freehep.org/>
<http://java.freehep.org/jaida>
<http://java.freehep.org/aidajni>
<http://jas.freehep.org/jas3>

Wide Area Networking / 446

ALICE Multi-site Data Transfer Tests on a Wide Area Network

A. Fritz¹ ; D. DiBari² ; D. Mura³ ; E. Fragiaco⁴ ; F. Carminati⁵ ; F. Minafra⁶ ; G. Donvito⁶ ; G. Lo Re⁷ ; J. Svec⁸ ; M. Masera⁹ ; M. Sitta¹⁰ ; P. Cerello¹⁰ ; R. Barbera¹¹ ; R. Turrisi¹² ; S. Bagnasco¹⁰

¹ *University of Houston*

² *Bari University*

³ *INFN Cagliari*

⁴ *INFN Trieste*

⁵ *CERN*

⁶ *INFN Bari*

⁷ *INFN & CNAF Bologna*

⁸ *ASCR Prague*

⁹ *Torino University*

¹⁰ *INFN Torino*

¹¹ *Catania University and INFN*

¹² *INFN Padova*

Auteur(s) contact: giuseppe.lore@cnaif.infn.it, federico.carminati@cern.ch

Next generation high energy physics experiments planned at the CERN Large Hadron Collider is so demanding in terms of both computing power and mass storage that data and CPU's can not be concentrated in a single site and will be distributed on a computational Grid according to a "multi-tier".

LHC experiments are made of several thousands of people from a few hundreds of institutes spread out all over the world. These people, according to their collaborations on specific physics analysis topics, can constitute highly dynamic Virtual Organizations rapidly changing as a function of both time and topology. The impact of future experiments on Wide Area Networks (WAN) will be non negligible especially for what concerns the capillarity of bandwidths (down to the "last mile"), quality of service, adaptivity and configurability.

In this paper we report on a series of multi-site data transfer tests performed within the ALICE Experiment on a wide area network test-bed in order to spot possible bottlenecks and pin down critical elements and parameters of actual research networks.

In order to make the tests as realistic as possible, reflecting the real use cases foreseen in the next future, we have taken into account all the aspects of the elements involved in the transfer of a file:

- Local disk Input/Output (I/O) performance;
- I/O block size;
- TCP parameters and number of parallel streams;
- Bandwidth Delay Product (BDP) expressed as the product of the Bandwidth (BW) times the Round Trip Time (RTT).

Distributed Computing Systems and Experiences / 48**AMS-02 Computing and Ground Data Handling**

A. Klimentov¹ ; M. Pohl² ; V. Choutko²

¹ *A*

² *E*

Auteur(s) contact: alexei.klimentov@cern.ch

AMS-02 Computing and Ground Data Handling.

V.Choutko (MIT, Cambridge), A.Klimentov (MIT, Cambridge) and M.Pohl (Geneva University)

AMS (Alpha Magnetic Spectrometer) is an experiment to search in space for dark matter and antimatter on the International Space Station (ISS). The AMS detector had a precursor flight in 1998 (STS-91, June 2-12, 1998). More than 100M events were collected and analyzed.

The final detector (AMS-02) will be installed on ISS in the fall of 2007 for at least 3 years. The data will be transmitted from ISS to NASA Marshall Space Flight Center (MSFC, Huntsville, Alabama) and transferred to CERN (Geneva Switzerland) for processing and analysis.

We are presenting the AMS-02 Ground Data Handling scenario and requirements to AMS ground centers: the Payload Operation and Control Center (POCC) and the Science Operation Center (SOC).

The Payload Operation and Control Center is where AMS operations take place, including commanding, storage and analysis of house keeping data and partial science data analysis for rapid quality control and feed back.

The AMS Science Data Center receives and stores all AMS science and house keeping data, as well as ancillary data from NASA. It ensures full science data reconstruction, calibration and alignment; it keeps data available for physics analysis and archives all data.

We also discuss the AMS-02 distributed MC production currently running in 15 Universities and Labs in Europe, USA and Asia, with automatic jobs submission and control from one central place (CERN). The software uses CORBA technology to control and monitor MC production and an ORACLE relational database, to keep catalogues, event description as well as production and monitoring information.

Poster Session 2 - Tableau: 61 / 308

ARDA Project Status Report

M. Lamanna¹ ; The ARDA Team²

¹ *CERN-IT*

² *CERN*

The ARDA project was started in April 2004 to support the four LHC experiments (ALICE, ATLAS, CMS and LHCb)

in the implementation of individual production and analysis environments based on the EGEE middleware.

The main goal of the project is to allow a fast feedback between the experiment and the middleware development teams via the construction and the usage of end-to-end prototypes allowing users to perform analyses out of the present data sets from recent montecarlo productions.

In this talk the project is presented with highlights of the first results and lessons learnt so far. The relations of the project with similar initiatives within and outside the High Energy Physics community are reviewed (notably in the EGEE application identification and support).

Distributed Computing Systems and Experiences / 503

ATLAS Data Challenge Production on Grid3

A. Vaniachine¹ ; E. May¹ ; G. Gieraltowski¹ ; H. Severini² ; K. De³ ; M. Mambelli⁴ ; M. Sosebee³ ; N. Ozturk³ ; P. McGuigan³ ; P. Nevski⁵ ; R. Baker⁶ ; R. Gardner⁴ ; W. Deng⁶ ; X. Zhao⁴ ; Y. Smirnov⁴

¹ ARGONNE NATIONAL LABORATORY

² Oklahoma University

³ University of Texas at Arlington

⁴ UNIVERSITY OF CHICAGO

⁵ BROOKHAVEN NATIONAL LABORATORY

⁶ Brookhaven National Laboratory

Auteur(s) contact: marco@hep.uchicago.edu, robert.w.gardner@cern.ch

We describe the design and operational experience of the ATLAS production system as implemented for execution on Grid3 resources. The execution environment consisted of a number of grid-based tools: Pacman for installation of VDT-based Grid3 services and ATLAS software releases, the Capone execution service built from the Chimera/Pegasus virtual data system for directed acyclic graph (DAG) generation, DAGMan/Condor-G for job submission and management, and the Windmill production supervisor which provides the messaging system for distributing production tasks to Capone. Produced datasets were registered into a distributed replica location service (Globus RLS) that was integrated with the Don Quixote proxy service for interoperability with other Grids used by ATLAS. We discuss performance, scalability, and fault handling during the first phase of ATLAS Data Challenge 2.

Poster Session 3 - Tableau: 61 / 212

ATLAS Detector Description Database Architecture

A. Vaniachine¹ ; D. Malon¹ ; K. Bernardet² ; P. Nevski³ ; S. Baranov⁴

¹ ANL

² CPPM

³ BROOKHAVEN NATIONAL LABORATORY

⁴ Freiburg University

Auteur(s) contact: vaniachine@anl.gov

In addition to the well-known challenges of computing and data handling at LHC scales, LHC experiments have also approached the scalability limit of manual management and control of the steering parameters ("primary numbers") provided to their software systems. The laborious task of detector description benefits from the implementation of a scalable relational database approach. We have created and extensively exercised in the ATLAS production environment a primary numbers database utilizing NOVA relational database technologies. In our report we describe the architecture of the relational database deployed for the storage, management, and uniform treatment of primary numbers in ATLAS detector description. We describe the benefits of the ATLAS software framework (Athena) on-demand data access architecture, and an automatic system for code generation of more than 300 classes (about 10% of ATLAS offline code) for primary numbers access from the Athena framework. Integration with the LHC Interval-of-Validity database infrastructure, measures for tighter primary numbers database input control, experience with ATLAS Combined Testbeam geometry and conditions payload storage using NOVA technologies integrated with the LHC ConditionsDB implementation, methods for application-side resource pooling, new user tools for knowledge discovery, navigation and browsing, and plans for new primary numbers database developments, are also described.

Distributed Computing Systems and Experiences / 115

ATLAS Distributed Analysis

A. Soroko¹ ; C. Haeberli² ; C. Kannan³ ; C.L. Tan⁴ ; D. Adams⁵ ; D. Liko⁶ ; F. Fassi⁷ ; F. Orellana⁶ ; G. Rybkine^{None} ; J. Fulachier⁸ ; J. Lozano⁷ ; K. Harrison⁹ ; M. Branco⁶ ; N. Chetan³ ; S. Albrand⁸ ; V. Sambamurthy³ ; W. Deng⁵

¹ *Oxford University*

² *Bern*

³ *Stonybrook*

⁴ *Birmingham*

⁵ *BNL*

⁶ *CERN*

⁷ *IFIC Valencia*

⁸ *LPSC, Grenoble*

⁹ *Cambridge University*

Auteur(s) contact: dladams@bnl.gov

The ATLAS distributed analysis (ADA) system is described. The ATLAS experiment has more than 2000 physicists from 150 institutions in 34 countries. Users, data and processing are distributed over these sites. ADA makes use of a collection of high-level web services whose interfaces are expressed in terms of AJDL (abstract job definition language) which includes descriptions of datasets, transformations and jobs. The high-level services are implemented using generic parts of these objects while clients and endpoint applications additionally make use of experiment-specific extensions. The key high-level service is the analysis service which receives a generic job request and creates and runs a corresponding job, typically as a collection of sub-jobs each handling a subset of the input dataset. The submitting client is able to monitor the progress of the job including partial results. The system is capable of running a wide range of applications but the emphasis is on event processing, in particular simulation, reconstruction and analysis of ATLAS data. Other high-level services include catalogs and dataset splitters and mergers. The ATLAS production system has been used to construct an analysis service

that makes production activities available to ATLAS users. An analysis service with interactive response is provided by DIAL. Another analysis service based on the EGEE middleware is being constructed in the context of the ARDA project. All are accessible from ROOT and python command lines and from the user-friendly graphical interface provided by GANGA.

Core Software / 134

ATLAS Metadata Interfaces (AMI) and ATLAS Metadata Catalogs

J. Fulachier¹ ; S. Albrand¹

¹ *LPSC*

Auteur(s) contact: solveig.albrand@lpsc.in2p3.fr

The ATLAS Metadata Interface (AMI) project provides a set of generic tools for managing database applications. AMI has a three-tier architecture with a core that supports a connection to any RDBMS using JDBC and SQL. The middle layer assumes that the databases have an AMI compliant self-describing structure. It provides a generic web interface and a generic command line interface. The top layer contains application specific features. The principal uses of AMI are the ATLAS Data Challenge dataset bookkeeping catalogs, and Tag Collector, a tool for release management.

The first AMI Web service client was introduced in early 2004. It offers many advantages over earlier clients because:

- Web services permit multi-language and multi-operating system support
- The user interface is very effectively de-coupled from the implementation.

Most upgrades can be implemented on the server side; no redistribution of client software is needed. In 2004 this client will be used for the ATLAS Data Challenge 2, for the ATLAS combined test beam offline bookkeeping, and also in the first prototypes of ARDA compliant analysis interfaces.

Distributed Computing Systems and Experiences / 501

ATLAS Production System in ATLAS Data Challenge 2

K. DE¹ ; L. GOOSSENS²

¹ *UNIVERSITY OF TEXAS AT ARLINGTON*

² *CERN*

Auteur(s) contact: luc.goossens@cern.ch

In order to validate the Offline Computing Model and the complete software suite, ATLAS is running a series of Data Challenges (DC). The main goals of DC1 (July 2002 to April 2003) were the preparation and the deployment of the

software required for the production of large event samples, and the production of those samples as a worldwide distributed activity.

DC2 (May 2004 until October 2004) is divided into three phases: (i) Monte Carlo data are produced using GEANT4 on three different Grids, LCG, Grid3 and NorduGrid; (ii) simulate the first pass reconstruction of data expected in 2007, also called Tier0 exercise, using the MC sample; and (iii) test the Distributed Analysis model.

A new automated data production system has been developed for DC2. The major design objectives are minimal human involvement, maximal robustness, and interoperability with several grid flavors and legacy systems. A central component of the production system is the production database holding information about all jobs. Multiple instances of a 'supervisor' component pick up unprocessed jobs from this database, distribute them to 'executor' processes, and verify them after execution. The 'executor' components interface to a particular grid or legacy flavour. The job distribution model is a combination of push and pull. A data management system keeps track of all produced data and allows for file transfers.

The basic elements of the production system are described. Experience with the use of the system in world-wide DC2 production of ten million events will be presented. We also present how the three Grid flavors are operated and monitored. Finally we discuss the first attempts on using the Distributed Analysis system.

Poster Session 3 - Tableau: 62 / 59

Adaptive Multi-vertex fitting

R. Fruehwirth¹ ; W. Waltenberger²

¹ *HEPHY Vienna*

² *HEPHY VIENNA*

Auteur(s) contact: walten@cern.ch

State of the art in the field of fitting particle tracks to one vertex is the Kalman technique. This least-squares (LS) estimator is known to be ideal in the case of perfect assignment of tracks to vertices and perfectly known Gaussian errors. Experimental data and detailed simulations always depart from this perfect model. The imperfections can be expected to be larger in high luminosity experiments like at the LHC. In such a context vertex fitting algorithms will have to be able to deal with mis-associated tracks and mis-estimated or non-Gaussian track errors. We present a vertex fitting technique that is insensitive to outlying observations and mis-estimated track errors, while it retains close-to-optimal results in the case of perfect data; it adapts to the data. This is realized by introducing weights that are associated to the tracks and reflect the compatibility of the tracks with the vertex. Outliers are no longer simply discarded - as is done in most of the classical robustification techniques - but rather downweighted. The algorithm will be presented in detail, and

comparisons with classical methods will be shown in relevant physics cases.

Event Processing / 259

Adding Kaons to the Bertini Cascade Model

A. Heikkinen¹ ; D. Wright²

¹ *Helsinki Institute of Physics HIP*

² *SLAC*

A version of the Bertini cascade model for hadronic interactions is part of the Geant4 toolkit, and may be used to simulate pion-, proton-, and neutron-induced reactions in nuclei. It is typically valid for incident energies of 10 GeV and below, making it especially useful for the simulation of hadronic calorimeters. In order to generate the intra-nuclear cascade, the code depends on tabulations of exclusive channel cross section data, parameterized angular distributions and phase-space generation of multi-particle final states. To provide a more detailed treatment of hadronic calorimetry, and kaon interactions in general, this model is being extended to include incident kaons up to an energy of 15 GeV. Exclusive channel cross sections, up to and including six-body final states, will be included for K⁺, K⁻, K₀, K₀bar, lambda, sigma⁺, sigma₀, sigma⁻, xi₀ and xi⁻. K⁺nucleon and K-nucleon cross sections are taken from various cross section catalogs, while most of the cross sections for incident K₀, K₀bar and hyperons are estimated from isospin and strangeness considerations. Because there is little data for incident hyperon cross sections, use of the extended model will be restricted to incident K⁺, K⁻, K₀S and K₀L. Hyperon cross sections are included only to handle the secondary interactions of hyperons created in the intra-nuclear cascade.

Core Software / 262

Addressing the persistency patterns of the time evolving HEP data in the ATLAS/LCG MySQL Conditions Databases

A. Amorim¹ ; A. Valassi² ; D. Duellmann² ; D. Klose¹ ; L. Pedro¹ ; N. Barros¹ ; T. Franco¹

¹ *FACULTY OF SCIENCES OF THE UNIVERSITY OF LISBON*

² *CERN*

Auteur(s) contact: antonio.amorim@fisica.fc.ul.pt, luis.pedro@cern.ch

The size and complexity of the present HEP experiments represents an enormous effort in the persistency of data. These efforts imply a tremendous investment in the databases field not only for the event data but also for data that is needed to qualify this one - the Conditions Data.

In the present document we'll describe the strategy for addressing the Conditions data problem in the ATLAS experiment, focusing in the ConditionsDB MySQL for the ATLAS/LCG project. The need for a persistent engine for structured conditions data has motivated the studies for an relational backend that maps transient structured objects in the relational database persistent engine. This paper illustrate the proposal for the storage of Conditions data

in the LCG framework using it both to store only the Interval Of Validity (IOV) and a reference that represents the 'path' to an external persistent storage mechanism, and to archive the IOV and the data in relational tables mapping the customizable CondDBTable objects. This allow to take advantages of all the relational features and also to directly map between transient objects and tables in the database server.

The issue of distributed data storage and partitioning, is also analyzed in this paper, taking into account the different levels of indirection that are provided by the ConditionsDB MySQL implementation. These features represent a very important built in functionality in terms of scalability, data balance in a system that aims to be completely distributed and with a very high performance for hundreds of users.

Distributed Computing Systems and Experiences / 443

AliEn Web Portal

A. Peters¹ ; L. Betev¹ ; M. Meoni² ; P E. Tissot-Daguette¹ ; P. Buncic¹ ; P. Saiz¹ ; V. Pinto Morais¹

¹ CERN

² CERN and INFN Firenze

Auteur(s) contact: predrag.buncic@cern.ch

The AliEn system, an implementation of the Grid paradigm developed by the ALICE Offline Project, is currently being used to produce and analyse Monte Carlo data at over 30 sites on four continents. The AliEn Web Portal is built around Open Source components with a backend based on Grid Services and compliant with the OGSA model. An easy and intuitive presentation layer gives the opportunity to the user to access information from multiple sources in a transparent and convenient way. Users can browse job provenance and access monitoring information from MonaLisa repository.

The presentation layer is separated from the content layer which is implemented via Grid and Web Services serving one or more users or Virtual Organizations.

Security and authentication of the portal are based on the Globus Grid Security infrastructure, OGSI::Lite and MyProxy online credentials repository.

In this presentation the architecture and functionality of AliEn Portal implementation will be presented.

Poster Session 2 - Tableau: 31 / 337

Alibaba: A heterogeneous grid-based job submission system used by the BaBar experiment

A. Forti¹ ; A. McNab² ; M. Jones¹ ; R. Barlow¹

¹ Manchester University

² UNIVERSITY OF MANCHESTER

Auteur(s) contact: mike.jones@man.ac.uk, roger.barlow@man.ac.uk

The BaBar experiment has accumulated many terabytes of data on particle physics reactions, accessed by a community of hundreds of

users.

Typical analysis tasks are C++ programs, individually written by the user, using shared templates and libraries. The resources have outgrown a single platform and a distributed computing model is needed. The grid provides the natural toolset. However, in contrast to the LHC experiments, BaBar has an existing user community with an existing non-Grid usage pattern, and providing users with an acceptable evolution presents a challenge.

The 'Alibaba' system, developed as part of the UK GridPP project, provides the user with a familiar command line environment. It draws on the existing global file systems employed and understood by the current user base. The main difference is that they submit jobs with a 'gsub' command that looks and feels like the familiar 'qsub'. However it enables them to submit jobs to computer systems at different institutions, with minimal requirements on the remote sites. Web based job monitoring is also provided. The problems and features (the input and output sandboxes, authentication, data location) and their solutions are described.

Online Computing / 61

An Embedded Linux System Based on PowerPC

C. YUANPING¹ ; M. Ye² ; T. KEJUN¹ ; Z. JINGWEI³

¹ IHEP, Beijing, China

² INSTITUTE OF HIGH ENERGY PHYSICS, ACADEMIA SINICA

³ IHEP, Beijing, CHINA

Auteur(s) contact: yem@ihep.ac.cn

This article introduces an Embedded Linux System based on vme series PowerPC as well as the base method on how to establish the system. The goal of the system is to build a test system of VMEbus device. It also can be used to setup the data acquisition and control system. Two types of compiler are provided by the developer system according to the features of the system and the PowerPC. At the top of the article some typical embedded Operation system will be introduced and the features of different system will be provided. And then the method on how to build an embedded Linux system as well as the key technique will be discussed in detail. Finally a successful data acquisition example will be given based on the test system.

Event Processing / 153

An Object-Oriented Simulation Program for CMS

A. De Roeck¹ ; A. FANFANI² ; A. NIKITENDO³ ; A. ZHOKIN⁴ ; D. ELVIRA⁵ ; F. FERRO⁶ ; G. EULISSE⁷ ; H.P. WELLISCH⁸ ; I. Osborne⁹ ; J. WENG¹⁰ ; K. LASSILA-PERINI¹¹ ; L. TUURA¹² ; M. CASE¹³ ; M. Liendl¹ ; M. Stavrianakou¹⁴ ; M. ZANETTI¹⁵ ; N. AMAPANE¹⁶ ; P. ARCE¹⁷ ; P. ZYCH¹⁸ ; S. ABDOULLINE¹⁹ ; S. BANERJEE²⁰ ; S. DUTTA²¹ ; S. MUZZAFFAR⁷ ; S. WYNHOFF²² ; T. BOCCALI²³ ; T. TODOROV²⁴ ; T. WILDISH²² ; V. ANDREEV²⁵

¹ CERN

² University of Bologna, Italy

- ³ *Imperial College, London, UK*
⁴ *ITEP, Moscow, Russia*
⁵ *FNAL, Chicago, USA*
⁶ *INFN, Genova, Italy*
⁷ *NorthEastern University, Boston, USA*
⁸ *CERN, Geneva, Switzerland*
⁹ *Northeastern University, Boston, USA*
¹⁰ *Karlsruhe University, Germany*
¹¹ *HIP, Helsinki, Finland*
¹² *NorthEastern Universty, Boston, USA*
¹³ *UC Davis, USA*
¹⁴ *FNAL*
¹⁵ *INFN, Italy*
¹⁶ *University of Torino, Italy*
¹⁷ *CIEMAT, Madrid, Spain*
¹⁸ *Warsaw University, Poland*
¹⁹ *FNAL, USA*
²⁰ *Tata Institute, Bombay, India*
²¹ *University of Pisa, Italy*
²² *Princeton University, USA*
²³ *SNS, Pisa, Italy*
²⁴ *IRES, Strasbourg, France*
²⁵ *UCLA, USA*

Auteur(s) contact: maya.stavrianakou@cern.ch

The CMS detector simulation package, OSCAR, is based on the Geant4 simulation toolkit and the CMS object-oriented framework for simulation and reconstruction.

Geant4 provides a rich set of physics processes describing in detail electro-magnetic and hadronic interactions. It also provides the tools for the implementation of the full CMS detector geometry and the interfaces required for recovering information from the particle tracking in the detectors.

This functionality is interfaced to the CMS framework, which, via its “action on demand” mechanisms, allows the user to selectively load desired modules and to configure and tune the final application.

The complete CMS detector is rather complex with more than 12 million readout channels and more than 1 million geometrical volumes.

OSCAR has been validated by comparing its results with test beam data and with results from simulation with a GEANT3-based program.

It has been successfully deployed in the 2004 data challenge for CMS, where ~20 million events for various LHC physics channels were simulated and analysed.

Authors:

S. Abdulline, V. Andreev, P. Arce, S. Arcelli, S. Banerjee, T. Boccali,
M. Case, A. De Roeck, S. Dutta, G. Eulisse, D. Elvira, A. Fanfani, F. Ferro,
M. Liendl, S. Muzaffar, A. Nikitenko, K. Lassila-Perini, I. Osborne,
M. Stavrianakou, T. Todorov, L. Tuura, H.P. Wellisch, T. Wildish, S. Wynnhoff,
M. Zanetti, A. Zhokin, P. Zych

Poster Session 2 - Tableau: 18 / 327

An intelligent resource selection system based on neural network for optimal application performance in a grid environment

G. PISCITELLI^{None} ; M. Castellano¹ ; T. Coviello¹

¹ DEE – POLITECNICO DI BARI, V. ORABONA, 4, 70125 – BARI, ITALY

Auteur(s) contact: castellano@deemail.poliba.it

Grid computing is a large scale geographically distributed and heterogeneous system that provides a common platform for running different grid enabled applications. As each application has different characteristics and requirements, it is a difficult task to develop a scheduling strategy able to achieve optimal performance because application-specific and dynamic system status have to be taken into account.

Moreover it may be possible to obtain optimal performance for multiple application simultaneously using a single scheduler. Hence in a lot of cases the application scheduling strategy is assigned to an expert application user who provides a ranking criterion for selecting the best computational element on a set of available resources. Such criteria are based on user perception of system capabilities and knowledge about the features and requirements of his application.

In this paper an intelligent mechanism has been both implemented and evaluated to select the best computational resource in a grid environment from the application viewpoint.

A neural network based system has been used to capture automatically the knowledge of a grid application expert user. The system scalability problem is also tackled and a preliminary solution based on sorting algorithm is discussed. The aim is to allow a common grid application user to benefit of this expertise.

Distributed Computing Systems and Experiences / 455

Application of the SAMGrid Test-Harness for Performance Evaluation and Tuning of a Distributed Cluster Implementation of Data Handling Services

A. Baranovski¹ ; A. Kreymer¹ ; A. Lyon¹ ; A. Sill² ; F. Ratnikov³ ; G. Garzoglio¹ ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter⁴ ; L. Lueking¹ ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber¹ ; R. Illingworth¹ ; R. Kennedy¹ ; R. St.Denis⁵ ; S. Belforte⁷ ; S. Stonjek⁸ ; S. White¹ ; U. Kerzel⁹ ; V. Bartsch⁶ ; W. Merritt¹

¹ FERMI NATIONAL ACCELERATOR LABORATORY

² Texas Tech University

³ Rutgers University

⁴ Fermilab

⁵ Glasgow University

⁶ Oxford University

⁷ INFN/Trieste

⁸ FERMI NATIONAL ACCELERATOR LABORATORY/Oxford University

⁹ Karlsruhe University

Auteur(s) contact: lyon@fnal.gov

The SAMGrid team has recently refactored its test harness suite for greater flexibility and easier configuration. This makes possible more interesting applications of the test harness, for component tests, integration tests, and stress tests. We report on the architecture of the test harness and its recent application to stress tests of a new analysis cluster at Fermilab, to explore the extremes of analysis use cases and the relevant parameters for tuning in the SAMGrid station services. This reimplementation of the test

harness is a python framework which uses XML for configuration and small plug-in python modules for specific test purposes. One current testing application is running on a 128-CPU analysis cluster with access to 6 TB distributed cache and also to a 2 TB centralized cache, permitting studies of different cache strategies. We have studied the service parameters which affect the performance of retrieving data from tape storage as well. The use cases studied vary from those which will require rapid file delivery with short processing time per file, to the opposite extreme of long processing time per file. We also show how the same harness can be used to run regular unit tests on a production system to aid early fault detection and diagnosis. These results are interesting for their implications with regard to Grid operations, and illustrate the type of monitoring and test facilities required to accomplish such performance tuning.

Event Processing / 103

Applications of the FLUKA Monte Carlo code in High Energy and Accelerator Physics

A. Empl¹ ; A. Fasso² ; A. Ferrari³ ; A. Ottolenghi⁴ ; D. Scannicchio⁵ ; E. Gadioli⁶ ; F. Ballarini⁷ ; F. Cerutti⁸ ; G. Battistoni⁹ ; J. Ranft¹⁰ ; K. TSOULOU¹¹ ; L. Pinsky¹² ; M V. Garzelli¹³ ; M. PELLICIONI¹⁴ ; P. Sala¹⁵ ; S. MURARO⁹ ; S. Roesler¹⁶ ; T. RANCATI⁹ ; V. PARINI¹⁷ ; V. Vlachoudis³

¹ *University of Houston, Texas, USA*

² *SLAC, Stanford, USA*

³ *CERN, Geneva, Switzerland*

⁴ *Universita' and INFN Pavia*

⁵ *Univeristy and INFN, Pavia, Italy*

⁶ *University and INFN, Milano, Italy*

⁷ *University and INFN Pavia, Italy*

⁸ *Univeristy and INFN Milano, Italy*

⁹ *INFN Milano, Italy*

¹⁰ *University of Siegen, Germany*

¹¹ *CERN Geneva, Switzerland*

¹² *University of Houston, Texas, USA*

¹³ *Univeristy and INFN, Milano, Italy*

¹⁴ *INFN, Frascati, Italy*

¹⁵ *ETH, Zurich, Switzerland*

¹⁶ *CERN, Geneva, Switzerland*

¹⁷ *INFN and University di Pavia, Italy*

Auteur(s) contact: giuseppe.battistoni@mi.infn.it

The FLUKA Monte Carlo transport code is being used for different applications in High Energy, Cosmic Ray and Accelerator Physics. Here we review some of the ongoing projects which are based on this simulation tool.

In particular, as far as accelerator physics is concerned, we wish to summarize the work in progress for the LHC and the CNGS project. From the point of view of experimental activity, a part the activity going in the framework of LHC detectors, we wish to discuss as a major example the application of FLUKA to the ICARUS Liquid Argon TPC.

Upgrades in cosmic ray calculations, to demonstrate the capability of FLUKA to reproduce existing experimental data, are also presented

Core Software / 87

Aspect-Oriented Extensions to HEP Frameworks

C. Tull¹ ; P. Calafiura¹

¹ *LBNL/ATLAS*

Auteur(s) contact: cetull@lbl.gov

In this paper we will discuss how Aspect-Oriented Programming (AOP) can be used to implement and extend the functionality of HEP architectures in areas such as performance monitoring, constraint checking, debugging and memory management. AOP is the latest evolution in the line of technology for functional decomposition which includes Structured Programming (SP) and Object-Oriented Programming (OOP). In AOP, an Aspect can contribute to the implementation of a number of procedures and objects and is used to capture a concern such as logging, memory allocation or thread synchronization that crosscuts multiple modules and/or types. We have chosen Gaudi as a representative HEP architecture because it is a component architecture and has been successfully adopted by several HEP experiments. Since most HEP frameworks are currently implemented in C++, for our study we have used AspectC++, an extension to C++ that allows the use of AOP techniques without adversely affecting software performance. We integrated AspectC++ in the development environment of the Atlas experiment, and we will discuss some of the configuration management issues that may arise in a mixed C++/AspectC++ environment. In this study we have focused on “Development Aspects”, i.e. aspects that are intended to facilitate program development but can be transparently removed from the production code, such as execution tracing, constraint checking and object lifetime monitoring. We will briefly discuss possible “Production Aspects” related to cache management and object creation. For each of the concerns we have examined we will discuss how traditional SP or OOP techniques compare to the AOP solution we developed. We will conclude discussing the short and medium term feasibility of introducing AOP, and AspectC++ in particular, in the complex software systems of the LHC experiments.

Core Software / 426

Aspects

J. Hrivnac¹

¹ *LAL*

Auteur(s) contact: julius.hrivnac@cern.ch

Aspect-Oriented Programming (AOP) is a new paradigm promising to allow further modularization of large software frameworks, like those developed in HEP. Such frameworks often manifest several orthogonal axes of contracts (Crosscutting Concerns - CC) leading to complex multidependencies. Currently used programming languages and development methodologies don't allow to easily identify and encapsulate such CC. AOP offers ways to solve CC problems by identifying places where they appear (Joint Points) and specifying actions to be applied at those places (Advices). While Aspects can be added in principle to any programming paradigm, they are mostly used in Object-Oriented environments. Thanks to wide acceptance and rich object model, most

Aspect-Oriented toolkits have been developed for Java language. Probably the most used AOP language is AspectJ.
The presentation will demonstrate using AspectJ language to solve several common HEP Crosscutting Concerns from simple cases (like logging or debugging) to complex ones (like object persistency, data analysis or graphics).

Grid Security / 104

Authentication/Security services in the ROOT framework

G. GANIS¹

¹ *CERN*

Auteur(s) contact: gerardo.ganis@cern.ch

The new authentication and security services available in the ROOT framework for client/server applications will be described.

The authentication scheme has been designed with the purpose to make the system complete and flexible, to fit the needs of the coming clusters and facilities.

Three authentication methods have been made available: Globus/GSI, for GRID-awareness; SSH, to allow using a secure and very popular protocol; a fast identification method for intrinsically secure situations.

A mechanism to allow server access control has been implemented, allowing to model the authorization schemes according to the needs.

A lightweight mechanism for client/server method negotiation has been introduced, to adapt to heterogeneous situations.

The forward of the authentication credentials in the PROOF system has been fully automatized.

The modularity of the code has been improved to ease maintenance and reuse in new ROOT modules. In particular, a plug-in library for the new Xrootd file server daemon has been designed and implemented.

Authentication support has been extended to the main socket server class, allowing to run a ROOT interactive session as a full-featured daemon.

Security services have also been added to ROOT. The exchange of sensitive information, e.g. passwords, has been secured.

New socket classes supporting SSL-secured connections have been provided for encryption of all the information exchanged with the remote host.

Poster Session 1 - Tableau: 1 / 118

AutoBlocker: A system for detecting and blocking of network scanning based on analysis of netflow data.

A. Bobyshev¹ ; D. Lamore¹ ; P. Demar¹

¹ *FERMILAB*

Auteur(s) contact: bobyshev@fnal.gov

In a large campus network, such as Fermilab's ten thousand nodes, scanning initiated from either outside of or within the campus network raises security concerns, may have very serious impact on network performance, and even disrupt normal operation of

many services. In this paper we introduce a system for detecting and automatic blocking of excessive traffic of different nature, scanning, DoS attacks, virus infected computers. The system, called AutoBlocker, is a distributed computing system based on quasi-real time analysis of network flow data collected from the border router and core routers. AutoBlocker also has an interface to accept alerts from the IDS systems (e.g. BRO, SNORT) that are based on other technologies. The system has multiple configurable alert levels for the detection of anomalous behavior and configurable trigger criteria for automated blocking of the scans at the core or border routers. It has been in use at Fermilab for about 2 years, and become a very valuable tool to curtail scan activity within the Fermilab campus network.

Poster Session 3 - Tableau: 59 / 54

Automated Tests in NICOS Nightly Control System

A. Undrus¹

¹ *BROOKHAVEN NATIONAL LABORATORY, USA*

Auteur(s) contact: undrus@bnl.gov

Software testing is a difficult, time-consuming process that requires technical sophistication and proper planning. This is especially true for the large-scale software projects of High Energy Physics where constant modifications and enhancements are typical. The automated nightly testing is the important component of NICOS, Nightly Control System, that manages the multi-platform nightly builds based on the recent versions of software packages. It facilitates collective work in collaborative environment and provides four benefits to developers: repeatability (tests can be executed more than once), accumulation (results are stored and reflected on NICOS web pages), feedback (automatic e-mail notifications about test failures), user friendly setup (configuration parameters can be encrypted in the body of test scripts). The modular structure of NICOS allows plugging in other validation and organization tools, such as QMTest and CppUNIT. NICOS classifies tests according to their granularity level and purpose. The low level structural tests reveal compilation problems, inconsistencies in package configuration, such as circular dependencies, and simple isolated bugs. The results for these three groups of tests are published for each package of the software project. The integrated (or behavioral) tests find bugs at levels of users scenarios and NICOS generates the special web page with their results. The NICOS tool is currently used to coordinate the efforts of more than 100 developers for the ATLAS project at CERN and included in the tool library of the LHC computing proje

Core Software / 291

Automatic Procedures as Generated Analysis Tool

G. Asova¹ ; K. Abrahamyan¹ ; M. Winde¹ ; P. Castro-Garcia² ; S. Weisse¹

¹ *DESY ZEUTHEN*

² DESY**Auteur(s) contact:** winde@ifh.de

The photo injector test facility at DESY Zeuthen (PITZ) was built to develop, operate and optimize photo injectors for future free electron lasers and linear colliders. In PITZ we use a DAQ system that stores data as a collection of ROOT files, forming our database for offline analysis. Consequently, the offline analysis will be performed by a ROOT application, written at least partly by the user (a physicist). To help the user to develop safe filters and data visualisation (graphs, histograms) with minimal effort in an existing ROOT framework application, we provide a GUI that generates C++ source files, compiles and links them to the rest of the application. We call these C++ routines “Automatic Procedures” (AP). Standard filter conditions and data visualisation can be generated by click or drag- and-drop, while more complex tasks may be expressed as small pieces of C++ code. Once compiled by ACLiC (ROOTs Automatic Compiler Linker), an Automatic Procedure may be reused without repeated compilation. E. g. the injector shift crew will run a number of ROOT applications, controlled by APs in regular intervals. Alternatively every AP can be read in and loaded to the GUI for further improvement. A number of APs can run in a logical sequence, parameters can be transferred from one AP to another. They can be selected by picking a point from a graph. The GUI was constructed with Qt, because that offers a comprehensive GUI programming toolkit.

The photo injector test facility at DESY Zeuthen (PITZ) was built to develop, operate and optimize photo injectors for future free electron lasers and linear colliders. In PITZ we use a DAQ system that stores data as a collection of ROOT files, forming our database for offline analysis. Consequently, the offline analysis will be performed by a ROOT application, written at least partly by the user (a physicist). To help the user to develop safe filters and data visualisation (graphs, histograms) with minimal effort in an existing ROOT framework application, we provide a GUI that generates C++ source files, compiles and links them to the rest of the application. We call these C++ routines “Automatic Procedures” (AP). Standard filter conditions and data visualisation can be generated by click or drag- and-drop, while more complex tasks may be expressed as small pieces of C++ code. Once compiled by ACLiC (ROOTs Automatic Compiler Linker), an Automatic Procedure may be reused without repeated compilation. E. g. the injector shift crew will run a number of ROOT applications, controlled by APs in regular intervals. Alternatively every AP can be read in and loaded to the GUI for further improvement. A number of APs can run in a logical sequence, parameters can be transferred from one AP to another. They can be selected by picking a point from a graph. The GUI was constructed with Qt, because that offers a comprehensive GUI programming toolkit.

Keywords: Automatic Procedure, ROOT, ACLiC, Data Analysis, Data Visualisation, GUI, Qt

Poster Session 3 - Tableau: 54 / 112

BFD: A software management & deployment tool for mixed-language distributed projects

Auteur(s): M. Stoufer¹

Co-auteur(s) S. Patton ¹

¹ LAWRENCE BERKELEY NATIONAL LAB

Auteur(s) contact: mcstoufer@lbl.gov

As any software project grows in both its collaborative and mixed codebase nature, current tools like CVS and Maven start to sag under the pressure of complex sub-project dependencies and versioning. A developer-wide failure in mastery of these tools will inevitably lead to an unrecoverable instability of a project. Even keeping a single software project stable in a large collaborative environment has proved a difficult venture in which numerous home-spun and commercial tools have yet to fully succeed.

BFD looks to solve the problems inherent with large scale software projects that span multiple software mixed-language projects. This is accomplished two-fold. BFD extends the versioning methodology of CVS or Maven by enforcing a rich data type format for its version tags. BFD also improves on the naive build ideologies of the developers IDE by being able to resolve complex dependencies between non-related projects as well as knowing when incompatible dependencies cannot be resolved. The concept of the Meta project has also been introduced to allow projects to be grouped together in a logical manner. Thus allowing varying versions of said projects to be kept track of by an overarching framework.

Distributed Computing Services / 338

BaBar Bookkeeping - a distributed meta-data catalog of the BaBar event store.

A. Ceseracciu¹ ; A. Forti² ; D. Bukin³ ; D. Hutchcroft⁴ ; D. Kovalskyi⁵ ; D. Smith⁶ ; G. Dubois-Felsmann⁷ ; P. Jackson⁸ ; T. Adye⁹ ; W. Roethel¹⁰

¹ INFN SEZIONE DI PADOVA

² UNIVERSITY OF MANCHESTER

³ BUDKER

⁴ UNIVERSITY OF LIVERPOOL

⁵ UNIVERSITY OF MARYLAND

⁶ STANFORD LINEAR ACCELERATOR CENTER

⁷ Caltech

⁸ UNIVERSITY OF LONDON, ROYAL HOLLOWAY

⁹ RUTHERFORD APPLETON LABORATORY

¹⁰ UC, IRVINE

Auteur(s) contact: douglas@slac.stanford.edu

The BaBar experiment has migrated its event store from an objectivity based system to a system using ROOT-files, and along with this has developed a new bookkeeping design. This bookkeeping now combines data production, quality control, event store inventory, distribution of BaBar data to sites and user analysis in one central place, and is based on collections of data stored as ROOT-files. These collections are grouped into pre-determined datasets, which define subsets of BaBar data to be used in analysis. Datasets are updated automatically to contain at any times the most up-to-date BaBar data. Local mirrors of the bookkeeping database can be used with the data distribution

features to import collections and maintain local event stores containing subsets of the available BaBar data. The bookkeeping system is scalable and supports sites containing all available data and hundreds of users down to the single user with a laptop. Oracle and MySQL relational databases are supported in its use, and sites can choose which to support. Database mirrors in the bookkeeping system can be accessed over network, which allows to browse local inventories from remote sites. This book keeping system has been in active use in BaBar since early this year, and the scope of its use along with technologies developed to keep it working will be presented.

Plenary / 502

BaBar computing - From collisions to physics results

P. ELMER¹

¹ *Princeton University*

Auteur(s) contact: peter.elmer@cern.ch

The BaBar experiment at SLAC studies B-physics at the Upsilon(4S) resonance using the high-luminosity e+e- collider PEP-II at the Stanford Linear Accelerator Center (SLAC). Taking, processing and analyzing the very large data samples is a significant computing challenge.

This presentation will describe the entire BaBar computing chain and illustrate the solutions chosen as well as their evolution with the ever higher luminosity being delivered by PEP-II. This will include data acquisition and software triggering in a high availability, low-deadtime online environment, a prompt, automated calibration pass through the data SLAC and then the full reconstruction of the data that takes place at INFN-Padova within 24 hours. Monte Carlo production takes place in a highly automated fashion in 25+ sites. The resulting real and simulated data is distributed and made available at SLAC and other computing centers.

For analysis a much more sophisticated skimming pass has been introduced in the past year, along with a reworked eventstore. This allows 120 highly customized analysis-specific skims to be produced for direct use by the analysis groups. This skim data format is the same eventstore data as that produced directly by the data and monte carlo productions and can be handled and distributed in the same way.

The total data volume in BaBar is about 1.5PB.

Distributed Computing Systems and Experiences / 339

BaBar simulation production - A millennium of work in under a year

C. Bozzi¹ ; D. Smith² ; F. Blanc³

¹ *INFN SEZIONE DI FERRARA*

² *STANFORD LINEAR ACCELERATOR CENTER*

³ *UNIVERSITY OF COLORADO*

Auteur(s) contact: douglas@slac.stanford.edu

for the BaBar Computing Group.

The analysis of the BaBar experiment requires many times the measured data to be produced in simulation. This requirement has resulted in one of the largest distributed computing projects ever completed.

The latest round of simulation for BaBar started in early 2003, and completed in early 2004, and encompassed over 1 million jobs, and over 2.2 billion events. By the end of the production cycle over 2 dozen different computing centers and nearly 1.5 thousand cpus were in constant use in North America and Europe. The whole effort was managed from a central database at SLAC, with real-time updates of the status of all jobs. Utilities were developed to tie together production with many different batch systems, and with different needs for security.

The produced data was automatically transferred to SLAC for use and distribution to analysis sites. The system developed to manage this effort was a combination of web and database applications, and command line utilities. The technologies used to complete this effort along with its complete scope will be presented.

Event Processing / 445

Bayesian Approach for Combined Particle Identification in ALICE Experiment at LHC

I. Belikov¹ ; K. Safarik¹ ; M. Ivanov¹ ; P. Hristov¹ ; T. Kuhr¹

¹ CERN

Auteur(s) contact: jouri.belikov@cern.ch

One of the main features of the ALICE detector at LHC is the capability to identify particles in a very broad momentum range from 0.1 GeV/c up to 10 GeV/c. This can be achieved only by combining, within a common setup, several detecting systems that are efficient in some narrower and complementary momentum sub-ranges. The situation is further complicated by the amount of data to be processed (about 10^7 events with about 10^4 tracks in each). Thus, the particle identification (PID) procedure should satisfy the following requirements:

- 1) It should be as much as possible automatic.
- 2) It should be able to combine PID signals of different nature (e.g. dE/dx and TOF measurements).
- 3) When several detectors contribute to the PID, the procedure must profit from this situation by providing an improved PID.
- 4) When only some detectors identify a particle, the signals from the other detectors must not affect the combined PID.
- 5) It should take into account the fact that the PID depends, due to different track selection, on the kind of analysis.

In this report we will demonstrate how combining the single detector PID signals in the Bayesian way satisfies these requirements. We will also discuss how one can obtain the needed probability distribution functions and a priory probability from the experimental data. The approach has been implemented within the

ALICE offline framework, and the algorithm efficiency and PID contamination have been estimated using the ALICE simulation.

Poster Session 2 - Tableau: 39 / 226

Beyond Persistence: Developments and Directions in ATLAS Data Management

A. Schaffer¹ ; D. Malon²

¹ *LAL Orsay*

² *ANL*

As ATLAS begins validation of its computing model in 2004, requirements imposed upon ATLAS data management software move well beyond simple persistence, and beyond the “read a file, write a file” operational model that has sufficed for most simulation production. New functionality is required to support the ATLAS Tier 0 model, and to support deployment in a globally distributed environment in which the preponderance of computing resources—not only CPU cycles but data services as well—reside outside the host laboratory.

This paper takes an architectural perspective in describing new developments in ATLAS data management software, including the ATLAS event-level metadata system and related infrastructure, and the mediation services that allow one to distinguish writing from registration and selection from retrieval, in a manner that is consistent both for event data and for time-varying conditions. The ever-broader role of databases and catalogs, and issues related to the distributed deployment thereof, are also addressed.

Poster Session 1 - Tableau: 21 / 467

Boosting the data logging rates in run 4 of the PHENIX experiment

C. Pinkenburg¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Auteur(s) contact: pinkenburg@bnl.gov

With the improvements in CPU and disk speed over the past years, we were able to exceed the original design data logging rate of 40MB/s by a factor of 3 already for the Run 3 in 2002. For the Run 4 in 2003, we increased the raw disk logging capacity further to about 400MB/s.

Another major improvement was the implementation of compressed data logging. The PHENIX raw data, after application of the standard data reduction techniques, were found to be further compressible by utilities like *gzip* by almost a factor of 2, and we defined a PHENIX standard of a compressed raw data format. The buffers that make up a raw data file consist of buffers that would get compressed and the resulting smaller data volume written out to disk. For a long time, this proved to be much too slow to be usable in the DAQ, until we could shift the compression to the event builder machines and so distributed the load over many fast CPU's. We also selected a

different compression algorithm, LZO, which is about a factor of 4 faster than the “compress2” algorithm used internally in gzip. With the compression, the raw data volume shrinks to about 60% of the original size, boosting the original data rate before compression to more than 700MB/s.

We will present the techniques and architecture, and the impact this has had on the data taking in Run 4.

Wide Area Networking / 177

Breaking the 1 GByte/sec Barrier? High speed WAN data transfers for science

D. NAE¹ ; H. Newman¹ ; J. Bunn¹ ; S. Ravot¹ ; X. SU¹ ; Y. Xia¹

¹ *Caltech*

Auteur(s) contact: sylvain.ravot@cern.ch, julian.bunn@caltech.edu

In this paper we describe the current state of the art in equipment, software and methods for transferring large scientific datasets at high speed around the globe. We first present a short introductory history of the use of networking in HEP, some details on the evolution, current status and plans for the Caltech/CERN/DataTAG transAtlantic link, and a description of the topology and capabilities of the research networks between CERN and HEP institutes in the USA. We follow this with some detailed material on the hardware and software environments we have used in collaboration with international partners (including CERN and DataTAG) to break several Internet2 land speed records over the last couple of years. Finally we describe our recent developments in collaboration with Microsoft, Newisys, AMD, Cisco and other industrial partners, in which we are attempting to transfer HEP data files from disk servers at CERN via a 10Gbit network path to disk servers at Caltech's Center for Advanced Computing Research (a total distance of over 11,000 kilometres), at a rate exceeding 1 GByte per second. We describe some solutions being used to overcome networking and hardware performance issues. Whilst such transfers represent the bleeding edge of what is possible today, they are expected to be commonplace at the start of LHC operations in 2007.

Wide Area Networking / 450

Bringing High-Performance Networking to HEP Users

N. Pezzi¹ ; R. Hughes-Jones² ; S. Dallison² ; Y-T. Li¹

¹ *University College London*

² *THE UNIVERSITY OF MANCHESTER*

Auteur(s) contact: steved@hep.man.ac.uk

How do we get High Throughput data transport to real users? The MB-NG project is a major collaboration which brings together expertise from users, industry, equipment providers and leading edge e-science application developers. Major successes in the areas of Quality of Service (QoS) and managed bandwidth have provided a leading edge U.K. Diffserv enabled network running at 2.5 Gbit/s. One of the central aims of MB-NG is the investigation of high performance data transport mechanisms for Grid data transfer across heterogeneous networks.

New transport stacks implement sender side modifications to the TCP algorithm which enable increased bandwidth utilisation in long-delay high-bandwidth environments. This allows a single stream of a modified TCP stack to transmit at rates that would otherwise require multiple streams of standard RENO TCP. This paper reports on investigations of the performance of these TCP stacks and their use with data transfer applications such as GridFTP, BBFTP, BBCP and APACHE. End-host performance behaviour was also examined in order to determine effects of the Network Interface, PCI bus performance, and disk and RAID sub-systems.

In a Collaboration between the BaBar experiment and MB-NG we demonstrated high performance data transport using these new TCP/IP transport protocol stacks and QoS provisioning. We report on the benefits of this introduction of high speed networks and advanced TCP stacks together with various levels of QoS to the BaBar computing environment. The benefits achieved are contrasted with network behaviour and application performance using today's "production" network.

Grid Security / 439

Building Global HEP Systems on Kerberos

M. Crawford¹

¹ FERMILAB

Auteur(s) contact: matt.crawford@cern.ch

As an underpinning of AFS and Windows 2000, and as a formally proven security protocol in its own right, Kerberos is ubiquitous among HEP sites. Fermilab and users from other sites have taken advantage of this and built a diversity of distributed applications over Kerberos v5. We present several projects in which this security infrastructure has been leveraged to meet the requirements of far-flung collaborations. These range from straightforward "Kerberization" of applications such as database and batch services, to quick tricks like simulating a user-authenticated web service with AFS and the "file:" schema, to more complex systems. Examples of the latter include experiment control room operations and the Central Analysis Farm (CAF).

We present several use cases and their security models, and examine how they attempt to address some of the outstanding problems of secure distributed computing: delegation of the least necessary privilege; establishment of trust between a user and a remote processing facility; credentials for long-queued or long-running processes, and automated processes running without any user's instigation; security of remotely-stored credentials; and ability to scale to the numbers of sites, machines and users expected in the collaborations of the coming decade.

Poster Session 2 - Tableau: 47 / 52

Building the LCG: from middleware integration to production quality software

C. Osuna¹ ; D. Qing² ; D. Smith³ ; F. Chollet⁴ ; G. Grosdidier⁵ ; J-P. Baud³ ; L. Poncet⁶ ; M. Litmaath³ ; M. Serra⁷ ; P. Bettini⁷ ; Z. Sekera³

¹ *University of Granada (Spain)*

² *Academia Sinica (Taiwan)*

³ *CERN*

⁴ *LAPP-IN2P3*

⁵ *LAL-IN2P3*

⁶ *LAL-IN2p3*

⁷ *INFN-CERN*

Auteur(s) contact: louis.poncet@cern.ch

In the last few years grid software (middleware) has become available from various sources. However, there are no standards yet which allow for an easy integration of different services.

Moreover, middleware was produced by different projects with the main goal of developing new functionalities rather than production quality software.

In the context of the LHC Computing Grid project (LCG) an integration, testing and certification activity is ongoing which aims at producing a stable coherent set of services.

Here we report on the processes employed to produce the LCG middleware release and related activities, including the infrastructures used, the activities needed to integrate the various components and the certification process.

Our certification process consists of a continuous iterative cycle that also involves feedback from the LCG production system and input from the software providers.

The architecture of the LCG middleware is described, including additional components developed by LCG to improve scalability and performance.

Other associated activities include packaging for deployment, porting to different platforms, debugging and patching of the software.

Functionality and stress tests are performed via a large test-bed infrastructure that allows for benchmarking of different configurations.

We describe also the results of our tests and our experience collected during the building of the LCG infrastructure.

Computer Fabrics / 230

CASTOR: Operational issues and new Developments

B. Couturier¹ ; J-D. Durand¹ ; O. Barring¹ ; S. Ponce¹

¹ *CERN*

Auteur(s) contact: jean-damien.durand@cern.ch

The Cern Advanced STORAGE (CASTOR) system is a scalable high throughput hierarchical storage system developed at CERN. CASTOR was first deployed for full production use in 2001 and has expanded to now manage around two PetaBytes and almost 20 million files. CASTOR is a modular system, providing a distributed disk cache, a stager, and a back end tape archive, accessible via a global logical name-space.

This paper focuses on the operational issues of the system currently in production, and first experiences with the new CASTOR stager which has undergone a significant redesign in order to cope with the data handling challenges posed by the LHC, which will be commissioned in 2007.

The design target for the new stager was to scale to another order of magnitude above the current CASTOR, namely to be able to sustain peak

rates of the order of 1000 file open requests per second for a PetaByte disk pool. The new developments have been inspired by the problems which arose managing massive installations of commodity storage hardware. The farming of disk servers poses new challenges to the disk cache management: request scheduling; resource sharing and partitioning; automated configuration and monitoring; and fault tolerance of unreliable hardware

Management of the distributed component based CASTOR system across a large farm, provides an ideal example of the driving forces for the development of automated management suites. Quattor and Lemon frameworks naturally address CASTOR's operational requirements, and we will conclude by describing their deployment on the masstorage systems at CERN.

Core Software / 316

CDB - Distributed Conditions Database of BaBar Experiment

D. Brown¹ ; I. Gaponenko²

¹ *Lawrence Berkeley National Laboratory (LBNL)*

² *LAWRENCE BERKELEY NATIONAL LABORATORY*

Auteur(s) contact: gapon@slac.stanford.edu

A new, completely redesigned Condition/DB was deployed in BaBar in October 2002. It replaced the old database software used through the first three and half years of data taking.

The new software aims at performance and scalability limitations of the original database. However this major redesign brought in a new model of the metadata, brand new technology- and implementation- independent API, flexible configurability and extended functionality.

One of the greatest strength of new CDB is that it's been designed to be a distributed kind database from the ground up to facilitate propagation and exchange of conditions (calibrations, detector alignments, etc.) in the realm of the international HEP collaboration.

The first implementation of CDB uses Objectivity/DB as its underlying persistent technology. There is an ongoing study to understand how to implement CDB on top of other persistent technologies.

The talk will cover the whole spectrum of topics ranging from the basic conceptual model of the new database through the way CDB is currently exploited in BaBar to the directions of further developments.

Poster Session 2 - Tableau: 66 / 488

CERN Modular Physics Screensaver or Using spare CPU cycles of CERN's Desktop PCs

A. Wagner¹ ; E.M. McIntosh¹

¹ *CERN*

Auteur(s) contact: andreas.wagner@cern.ch, eric.mcintosh@cern.ch

CERN has about 5500 Desktop PCs. These computers offer a large pool of resources that can be used for physics calculations outside office hours.

The paper describes a project to make use of the spare CPU cycles of these PCs for LHC tracking studies. The client server application is implemented as a lightweight,

modular screensaver and a Web Application containing the physics job repository. The information exchange between client and server is done using the HTTP protocol. The design and implementation is presented together with results of performance and scalability studies. A typical LHC tracking study involves some 1500 jobs, each over 100,000 turns, requiring about 1 hour of CPU on a modern PC. A reliable and easy to use Linux interface to the CPSS Web application has been provided. It has been used for a production run of 15,000 jobs, using some 50 desktop Windows PCs, which uncovered a numerical incompatibility between Windows 2000 and XP. It is expected to make available up to two orders of magnitude more computing power for these studies at zero cost.

Computer Fabrics / 239

CERN's openlab for Datagrid applications

S. Jarp¹

¹ CERN

Auteur(s) contact: sverre.jarp@cern.ch

For the last 18 months CERN has collaborated closely with several industrial partners to evaluate, through the opencluster project, technology that may (and hopefully will) play a strong role in the future computing solutions, primarily for LHC but possibly also for other HEP computing environments. Unlike conventional field testing where solutions from industry are evaluated rather independently, the openlab principle is based on active collaboration between all partners, with the common goal of constructing a coherent system.

The talk will discuss our experience to date with the following hardware
- 64-bit computing (in our case represented by the Itanium processor). This will also

include the porting of applications and Grid software to 64 bits.

- Rack mounted servers
- The use of 10 Gbps Ethernet for both LAN and WAN connectivity
- An iSCSI-based Storage System that promises to scale to Petabyte dimensions
- The use of 10 Gbps Infiniband as a cluster interconnect

On the software side we will review our experience with the latest grid-enabled release of Oracle, the so-called release "10g".

The talk will review the results obtained so far, either in stand alone tests or as part of the larger LCG testbed, and it will describe the plans for the future in this three-year collaboration with industry.

Event Processing / 261

CHIPS based hadronization of quark-gluon strings

M. Kosov¹

¹ CERN

Auteur(s) contact: kosov@cern.ch

Quark-gluon strings are usually fragmented on the light cone in hadrons (PITHIA, JETSET) or in small hadronic clusters which decay in hadrons (HERWIG). In both cases the transverse momentum distribution is parameterized as an unknown function. In CHIPS the colliding hadrons stretch Pomeron ladders to each other and, when the Pomeron ladders meet

in the rapidity space, they create Quasmons (hadronic clusters bigger than Amati-Veneziano clusters of HERWIG). The Quasmon size and the corresponding transverse momentum distributions are tuned by the Drell-Yan $\mu^+\mu^-$ pairs. The final Quasmon fragmentation in CHIPS is tuned by the e^+e^- and proton-antiproton annihilation, which is already published.

Computer Fabrics / 476

CHOS, a method for concurrently supporting multiple operating system

C.L. Whitney¹ ; S. Canon¹

¹ NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Auteur(s) contact: canon@nersc.gov

Supporting multiple large collaborations on shared compute farms has typically resulted in divergent requirements from the users on the configuration of these farms. As the frameworks used by these collaborations are adapted to use Grids, this issue will likely have a significant impact on the effectiveness of Grids. To address these issues, a method was developed at Lawrence Berkeley National Lab and is being used in production on the PDSF cluster. This method, termed CHOS, uses a combination of a Linux kernel module, the change root system call, and several utilities to provide access to multiple Linux distributions and versions concurrently on a single system. This method will be presented, along with an explanation on how it is integrated into the login process, grid services, and batch scheduler systems. We will also describe how a distribution is installed and configured to run in this environment and explore some common problems that arise. Finally, we will relate our experience in deploying this framework on a production cluster used by several high energy and nuclear physics collaborations.

Core Software / 312

CLHEP Infrastructure Improvements

L. Garren¹

¹ F.N.A.L.

Auteur(s) contact: garren@fnal.gov

CLHEP is a set of HEP-specific foundation and utility classes such as random number generators, physics vectors, and particle data tables. Although CLHEP has traditionally been distributed as one large library, the user community has long wanted to build and use CLHEP packages separately.

With the release of CLHEP 1.9, CLHEP has been reorganized and enhanced to enable building and using CLHEP packages individually as well as collectively. The revised build strategy employs all the components of the standard autotools suite: automake, autoconf, and libtool. In combination with the reorganization, the use of these components makes

it easy not only to rebuild any single package (e.g., when that package changes), but also to add new packages.

This presentation will discuss the new CLHEP structure, illustrate the role and use of the autotools, and describe how other packages with similar organization can be seamlessly integrated with the CLHEP libraries.

Poster Session 3 - Tableau: 66 / 53

CMD-3 Project Offline Software Development

Auteur(s): A. Zaytsev¹

Co-auteur(s) E. Algaer² ; N. Stuly² ; S. Pirogov¹

¹ *BUDKER INSTITUTE OF NUCLEAR PHYSICS*

² *NOVOSIBIRSK STATE UNIVERSITY*

CMD-3 is the general purpose cryogenic magnetic detector for VEPP-2000 electron-positron collider, which is being commissioned at Budker Institute of Nuclear Physics (BINP, Novosibirsk, Russia). The main aspects of physical program of the experiment are study of known and search for new vector mesons, study of the $p\bar{p}$ and $n\bar{n}$ production cross sections in the vicinity of the threshold and search for exotic hadrons in the region of center-of-mass energy below 2 GeV. The essential upgrade of CMD-2 detector (designed for VEPP-2M collider at BINP) farm and distributed data storage management software is required to satisfy new detector needs and scheduled to perform in near future.

In this talk I will present the general design overview and status of implementation of CMD-3 offline software for reconstruction, visualization, data farm management and user interfaces. Software design standards for this project are object oriented programming techniques, C++ as a main language, Geant4 as an only simulation tool, Geant4 and GDML based detector geometry description, WIRED and HepRep based visualization, CLHEP library based primary generators and Linux as a main platform. The dedicated software development framework (Cmd3Fwk) was implemented in order to be the basic software integration solution and persistency manager. We also look forward to achieve high level of integration with ROOT framework and Geant4 toolkit.

Core Software / 235

CMS Detector Description: New Developments

Auteur(s): M. Case¹

Co-auteur(s) A.J. Muhammad² ; A.T.M. Aerts³ ; M. Liendl⁴

¹ *UNIVERSITY OF CALIFORNIA, DAVIS*

² *University of the West of England*

³ *Eindhoven University of Technology*

⁴ *CERN*

Auteur(s) contact: case@ucdhep.ucdavis.edu

The CMS Detector Description Database (DDD) consists of a C++ API and an XML based detector description language. DDD is used by the CMS simulation (OSCAR), reconstruction (ORCA), and visualization (IGUANA) as well by test beam software that

relies on those systems. The DDD is a sub-system within the COBRA framework of the CMS Core Software. Management of the XML is currently done using a separate Geometry project in CVS.

We give an overview of the DDD integration and report on recent developments concerning detector description in CMS software:

- The ability of client software to describe sub-detectors by providing an algorithm plug-in in C++ based on SEAL plug-in facilities. A typical algorithm plug-in makes use of the DDD API to describe detector properties. Through the API seamless access to data defined via the XML description language is ensured.
- An Oracle schema was recently developed and the database populated by a DDD application. The geometrical structure of the detector is seen as a skeleton to which conditions or configuration data can be attached.
- A C++ streaming mechanism to output the geometry as binary files was developed. This representation can be read into memory much more rapidly than the XML files can be parsed.

The DDD API shields clients from each of the possible input sources. Even the simultaneous use of several different input sources is possible through various configuration options in the framework COBRA.

Poster Session 3 - Tableau: 47 / 137

CMS Software Installation

A. Nowack¹ ; A. Sciaba² ; J. Weng³ ; K. Rabbertz⁴ ; M. Corvo⁵ ; S. MUZAFFAR⁶ ; S. Wynhoff⁷

¹ RWTH Aachen

² CERN

³ CERN & University of Karlsruhe

⁴ UNIVERSITY OF KARLSRUHE

⁵ INFN Padova

⁶ NorthEastern University, Boston, USA

⁷ Princeton University

Auteur(s) contact: klaus.rabbertz@cern.ch

For data analysis in an international collaboration it is important to have an efficient procedure to distribute, install and update the centrally maintained software. This is even more true when not only locally but also grid accessible resources are to be exploited. A practical solution will be presented that has been successfully employed for CMS software installations on systems ranging from physicists' notebooks up to LCG2 enabled clusters. It is based on perl for an automated production of rpm's and xcmsi, a tool written in perl and perl/Tk, to facilitate installing, updating and verifying our rpm packaged software.

Poster Session 3 - Tableau: 34 / 66

CMS Tracker Visualisation Tools

A. Regano¹ ; G. Zito² ; I. Osborne³ ; M.S. Mennea⁴

¹ *University Bari*

² *INFN Bari*

³ *Northeastern University, Boston, USA*

⁴ *UNIVERSITY & INFN BARI*

Auteur(s) contact: maria.santa.mennea@cern.ch

This document will review the design considerations, implementations and performance of the CMS Tracker Visualization tools. In view of the great complexity of this subdetector (more than 50 millions channels organized in 17000 modules each one of these being a complete detector), the standard CMS visualisation tools (IGUANA and IGUANACMS) that provide basic 3D capabilities and integration within CMS framework respectfully have been complemented with additional 2D graphics objects and a detailed object model of the tracker. Based on the experience acquired by using this software to debug and understand both hardware and software during the construction phase, we will propose possible future improvements to cope with online monitoring and event analysis during data taking.

Poster Session 3 - Tableau: 38 / 99

Carrot ROOT Apache Module

F. Rademakers¹ ; R. Brun¹ ; V. Onuchin²

¹ *CERN*

² *CERN, IHEP*

Auteur(s) contact: valeri.onouchine@cern.ch

Carrot is a scripting module for the Apache webserver. Based on the ROOT framework, it has a number of powerful features, including the ability to embed C++ code into HTML pages, run interpreted and compiled C++ macros, send and execute C++ code on remote web servers, browse and analyse the remote data located in ROOT files with the web browser, access and manipulate databases, and generate graphics on-the-fly, among many others.

In this talk we will describe and demonstrate the main features of Carrot. We will also discuss the future development of this module in context of GRID and integration with PROOF and xrootd/rootd.

More information about Carrot is available from the Carrot website at: <http://carrot.cern.ch>

Poster Session 2 - Tableau: 65 / 33

Central Reconstruction System on the RHIC Linux Farm in Brookhaven Laboratory

A. Withers¹ ; C. Hollowell² ; G. Tsai¹ ; R. Hogue² ; R. Popescu¹ ; T. Chan³ ; T. Wlodek²

¹ *Brookhaven National Laboratory*

² *Brookhaven National Lab*

³ *BROOKHAVEN NATIONAL LAB*

Auteur(s) contact: tony@bnl.gov

A description of a Condor-based, Grid-aware batch software system configured to function asynchronously with a mass storage system is presented. The software is currently used in a large Linux Farm (2700+ processors) at the RHIC and ATLAS Tier 1 Computing Facility at Brookhaven Lab. Design, scalability, reliability, features and support issues with a complex Condor-based batch system are addressed within the context of a Grid-like, distributed computing environment.

Computer Fabrics / 343

Chimera - a new, fast, extensible and Grid enabled namespace service

M. Gasthuber¹ ; P. Fuhrmann¹ ; T. Mkrtchyan¹

¹ *DESY*

After successful implementation and deployment of the dCache system over the last years, one of the additional required services, the namespace service, is faced additional and completely new requirements. Most of these are caused by scaling the system, the integration with Grid services and the need for redundant (high availability) configurations. The existing system, having only an NFSv2 access path, is easy to understand and well accepted by the users. This single 'access path' limits data management task to make use of classical tools like 'find', 'ls' and others. This is intuitiv for most users, but failed while dealing with millions of entries (files) and more sophisticated organizational schemes (metadata). The new system should support a native programmable interface (deep coupled, but fast), the 'classical' NFS path (now version 3 at least),

a dCache native access and the SQL path allowing any type of metadata to be used in complex queries. Extensions with other 'access paths' will be possible. Based on the experience with the current system we highlight on the following requirements:

- large file support (64 Bit) + large number of files (> 10⁸)
- fast
- Platform independents (runtime + persistent objects)
- Grid name service integration
- custom dCache integration
- redundant, high available runtime configurations (concurrent backup etc.)
- user usable metadata (store and query)
- ACL support
- pluggable authentication (e.g. GSSAPI)
- external processes can register for namespace events (e.g. removal/creation of files)

The presentation will show a detailed analysis of the requirements, the choosen design and selection of existing components. The current schedule should allow to show the first prototype results.

Poster Session 1 - Tableau: 11 / 159

Cluster architectures used to provide CERN central CVS services

I. Reguero¹ ; M. Guijarro¹ ; S. Lopienski¹

¹ CERN

Auteur(s) contact: sebastian.lopienski@cern.ch

There are two cluster architecture approaches used at CERN to provide central CVS services. The first one (<http://cern.ch/cvs>) depends on AFS for central storage of repositories and offers automatic load-balancing and fail-over mechanisms.

The second one (<http://cern.ch/lcgcvs>) is an N + 1 cluster based on local file systems, using data replication and not relying on AFS. It does not provide either dynamic load-balancing or automatic fail-over. Instead a series of tools were developed for repository relocation in case of fail-over and for manual load-balancing.

Both architectures are used in production at CERN and project managers can chose one or the other, depending on their needs. If, eventually, one architecture proves to be significantly better, the other one may be phased out. This paper presents in detail both approaches and describes their relative advantages and drawbacks, as well as some data about them (number of repositories, average repository size, etc).

Poster Session 3 - Tableau: 43 / 336

Comparative study of the power of goodness-of-fit algorithms

A. Pfeiffer¹ ; A. Ribon¹ ; B. Mascialino² ; M.G. Pia³ ; S. Donadio² ; S. Guatelli⁴

¹ CERN

² INFN Genova

³ INFN GENOVA

⁴ INFN Genova, Italy

Auteur(s) contact: maria.grazia.pia@cern.ch

A Toolkit for Statistical Data Analysis has been recently released. Thanks to this novel software system, for the first time an ample set of sophisticated algorithms for the comparison of data distributions (goodness of fit tests) is made available to the High Energy Physics community in an open source product. The statistical algorithms implemented belong to two sets, for the comparison of binned and unbinned distributions respectively; they include the Chi-squared Test, the Kolmogorov-Smirnov Test, the Kuiper Test, the Goodman Test, the Anderson-Darling Test, the Fisz-Cramer-von Mises test, the Tiku Test.

Since the Toolkit provides the user a wide choice of algorithms, it is important to evaluate them comparatively and to estimate their power, to provide guidance to the users about the selection of the most appropriate algorithm for a given use case.

We present a study of the power of a variety of mathematical algorithms implemented in the Toolkit. The study is performed by evaluating the behaviour of the various tests in a set of well identified use cases relevant to data analysis applications. To our

knowledge, such a comparative study of the power of goodness of fit algorithms has never been performed previously.

Core Software / 64

Composite Framework for CMS Applications

G. Eulisse¹ ; I. Osborne² ; L. Taylor¹ ; L. Tuura¹ ; S. Muzaffar¹ ; V. Innocente³

¹ *Northeastern University*

² *Northeastern University, Boston, USA*

³ *CERN*

Auteur(s) contact: ianna.osborne@cern.ch

We present a composite framework which exploits the advantages of the CMS data model and uses a novel approach for building CMS simulation, reconstruction, visualisation and future analysis applications. The framework exploits LCG SEAL and CMS COBRA plug-ins and extends the COBRA framework to pass communications between the GUI and event threads, using SEAL callbacks to navigate through the metadata and event data interactively in a distributed environment.

We give examples of current applications based on this framework, including CMS test-beams, geometry description debugging, GEANT4 simulation, event reconstruction, and the verification of reconstruction and higher level trigger algorithms.

Plenary / 532

Computing Models and Data Challenges of the LHC experiments

D. STICKLAND^{None}

The LHC experiments are undertaking various data-challenges in the run-up to completion of their computing models and the submission of the experiment and of the LHC Computing Grid (LCG), Technical Design Reports(TDR) in 2005. In this talk we summarize the current working models of the LHC Computing Models, identifying their similarities and differences. We summarize the results and status of the data challenges and identify critical areas still to be tested

Plenary / 492

Computing for Belle

N. KATAYAMA¹

¹ *KEK*

The Belle experiment operates at the KEKB accelerator, a high luminosity asymmetric energy $e^+ e^-$ machine. KEKB has achieved the world highest luminosity of $1.39 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Belle accumulates more than 1 million $B \bar{B}$ pairs in one good day. This corresponds to about 1.2 TB of raw data per day. The amount of the raw and processed data accumulated so far exceeds 1.4 PB. Belle's computing model has been a traditional one and very successful so far. The computing has been managed by minimal number of people using cost effective solutions. Looking at the future, KEKB/Belle plans to improve the luminosity to a few times $10^{35} \text{ cm}^{-2}\text{s}^{-1}$, 10 times as much as we obtain now. This presentation describes Belle's efficient computing operations, struggles to manage large amount of raw and physics data, and plans for Belle computing for Super KEKB/Belle.

Plenary / 505

Concepts and technologies used in contemporary DAQ systems

M. Purschke¹

¹ *Brookhaven National Laboratory*

Auteur(s) contact: purschke@bnl.gov

The concepts and technologies applied in data acquisition systems have changed dramatically over the past 15 years. Generic DAQ components and standards such as CAMAC and VME have largely been replaced by dedicated FPGA and ASIC boards, and dedicated real-time operation systems like OS9 or VxWorks have given way to Linux-based trigger processor and event building farms. We have also seen a shift from standard or proprietary bus systems used in event building to GigaBit networks and commodity components, such as PCs. With the advances in processing power, network throughput, and storage technologies, today's data rates in large experiments routinely reach hundreds of MegaBytes/s.

We will present examples of contemporary DAQ systems from different experiments, try to identify or categorize new approaches, and will compare the performance and throughput of existing DAQ systems with the projected data rates of the LHC experiments to see how close we have come to accomplish these goals. We will also try to look beyond the field of High-Energy Physics and see if there are trends and technologies out there which are worth keeping an eye on.

Poster Session 3 - Tableau: 25 / 264

Conditions Databases: the interfaces between the different ATLAS systems

A. Amorim¹; A. Kazarov²; A. Perus³; A. Schaffer³; D. Burckhart-Chromek⁴; D. KLOSE⁵; D. Liko⁴; D. Malon⁶; E. Badescu⁷; I. Soloviev²; J. Cook⁴; J. Flammer⁴; L. Mapelli⁴; L. Pedro¹; M. Caprini⁷; M. Dobson⁴; M. Mineev⁸; N. Barros¹; R. Hawkings⁴; R. Jones⁹; S. Kolos²; T. Franco¹; V. Kotov⁸; Y. Alexandrov¹⁰; Y. Ryabov¹¹

¹ *FACULTY OF SCIENCES OF THE UNIVERSITY OF LISBON*

² *Petersburg Nuclear Physics Institute (PNPI)*

³ *Laboratoire de l'Accélérateur Lineaire (LAL) (IN2P3) (LAL) Université de Paris-Sud (ParisXI)*

⁴ *CERN*

⁵ *Universidade de Lisboa, Portugal*

⁶ *ANL*

⁷ *Institute of Atomic Physics National Institute for Physics and Nuclear Engineering (IFIN-HH)*

⁸ *Joint Institute for Nuclear Research (JINR)*

⁹ *CERN, Geneva, Switzerland*

¹⁰ *P.N. Lebedev Institute of Physics (FIAN) Russian Academy of Sciences*

¹¹ *High Energy Physics Division (HEPD) Theory Division Petersburg Nuclear Physics Institute (PNPI)*

Auteur(s) contact: luis.pedro@cern.ch

Conditions Databases are beginning to be widely used in the ATLAS experiment. Conditions data are time-varying data describing the state of the detector used to reconstruct the event data. This includes all sorts of slowly evolving data like detector alignment, calibration, monitoring and data from Detector Control System (DCS).

In this paper we'll present the interfaces between the ConditionsDB and the DCS, Trigger and Data Acquisition (TDAQ) and offline control framework (Athena).

In the DCS case, a PVSS API Manager was developed based on the C++ interface for the ConditionsDB. The Manager links to a selection of datapoints and stores any value change in the ConditionsDB. The structure associated to each datapoint is mapped to a table that reflects this structure and is stored in the database.

The ConditionsDB Interface to the TDAQ (CDI) is a service provided by the Online Software that acts as an intermediary between TDAQ producers and consumers of conditions data. CDI provides the pathway to the ConditionsDB information regarding the present or past condition of the detector and trigger system as well as all the operational and monitoring data. It will provide the link between the Information Service (IS) and the ConditionsDB

Conditions database integration into the ATLAS Athena framework is also described, including connections to Athena's transient interval-of-validity management, conversion services to support conditions data I/O into Athena transient stores, and mechanisms by which the conditions database may be used for timestamp-mediated access to data stored in other technologies such as NOVA and POOL.

Plenary / 519

Conference Closeout

W. VON RUEDEN^{None}

Plenary / 518

Conference Conclusions

L. BAUERDICK¹

¹ *FNAL*

Online Computing / 302

Control in the ATLAS TDAQ System

A. AMORIM¹ ; A. Kazarov² ; D. BURCKHART-CHROMEK³ ; D. KLOSE¹ ; D. Liko⁴ ; E. BADESCU⁵ ; I. ALEXANDROV⁶ ; I. SOLOVIEV⁷ ; J. FLAMMER³ ; L. MAPELLI³ ; L. PEDRO¹ ; M. CAPRINI⁵ ; M. DOBSON³ ; M. Mineev⁸ ; N. FIUZA DE BARROS¹ ; R. JONES³ ; S. KOLOS¹ ; S. KOROBV⁶ ; V. KotoV⁸ ; Y. RYABOV⁷

¹ *Universidade de Lisboa, Portugal*

² *Petersburg Nuclear Physics Institute (PNPI)*

³ *CERN, Geneva, Switzerland*

⁴ *CERN*

⁵ *National Institute for Nuclear Physics and Engineering, Bucharest, Romania*

⁶ *JINR, Dubna, Russia*

⁷ *PNPI, St-Petersburg, Russia*

⁸ *Joint Institute for Nuclear Research (JINR)*

Auteur(s) contact: dietrich.liko@cern.ch

The unprecedented size and complexity of the ATLAS TDAQ system requires a comprehensive and flexible control system. Its role ranges from the so-called run-control, e.g. starting and stopping the datataking, to error handling and fault tolerance. It also includes intialisation and verification of the overall system. Following the traditional approach a hierachical system of customizable controllers has been proposed. For the final system all functionality would be therefore available in a distributed manner, with the possibility of local customisation.

After a technology survey the open source expert systemn CLIPS has been chosen as a basis for the implementation of the supervision and the verification system. The CLIPS interpreter has been extended to provide a general control framework. Other ATLAS Online software components have been integrated as plugins and provide the mechanism for configuration and communication. Several components have been implemented, that share this technology. The dynamic behaviour of the individual component is fully described by the rules, while the framework is based on a common implementation. During these year these components have been the subject of scalability tests up to the full system size. Encouraging results are presented and validate the technology choice.

Poster Session 1 - Tableau: 19 / 258

Control and state logging for the PHENIX DAQ System

E. desmond¹

¹ *BROOKHAVEN NATIONAL LAB*

Auteur(s) contact: desmond2@bnl.gov

The PHENIX DAQ system is managed by a control system responsible for the configuration and monitoring of the PHENIX detector hardware and readout software. At its core, the control system, called Runcontrol, is a process that manages the various components by way of a distributed architecture using CORBA. The control system, called Runcontrol, is a set of process that manages virtually all detector components through a distributed architecture base on CORBA. A key aspect of the distributed control system, the messaging

system, is the ability to access critical detector state information, and deliver it to operators and applications of the control system. The goal of the system is to concentrate all output messages of the distributed processes, which would normally end up in log files or on a terminal, in a central place. The messages may originate from or be received by applications running on any of the multiple platforms which are in use including Linux, Windows, Solaris, and VxWorks. Listener applications allow the DAQ operators to get a comprehensive overview of all messages they are interested in, and also allows scripts or other programs to take automated action in response to certain messages.

Messages are formatted to contain information about the source of the message, the message type, and its severity. Applications written to provide filtering of messages by the DAQ operators by type, severity and source will be presented.

We will discuss the mechanism underlying this system, present examples of the use, and discuss performance and reliability issues.

Poster Session 2 - Tableau: 36 / 385

Cross Experiment Workflow Management: The Runjob Project

D. Evans¹ ; G. Graham¹ ; I. Bertram² ; P. Love²

¹ *Computing Division, Fermi National Accelerator Laboratory*

² *Lancaster University*

Auteur(s) contact: peter.love@cern.ch, i.bertram@lancaster.ac.uk

Building on several years of success with the MCRunjob projects at DZero and CMS, the Fermilab sponsored joint Runjob project aims to provide a Workflow description language common to three experiments: DZero, CMS and CDF. This project will encapsulate the remote processing experiences of the three experiments in an extensible software architecture using web services as a communication medium. The core of the Runjob project will be the Shahkar software packages that provide services for describing jobs and targeting them at different execution environments. A common interface to multiple storage and compute grid elements will be provided, allowing the three experiments to share hardware resources in a transparent manner. Several tools provided by Shahkar are discussed including FileMetaBrokers, which provide a uniform way to handle files and metadata over a distributed cluster, the ShREEK runtime execution environment that allows executable jobs to provide a real time monitoring and control interface to any system, the scriptObject generic task encapsulation objects and XMLProcessor object persistency tool.

Computer Fabrics / 489

Current Status of Fabric Management at CERN

D. Front¹ ; D. Waldron² ; G. Cancio² ; H. Meinhard³ ; J. van Eldik² ; M. Siket² ; M. Stepniewski² ; P. Poznanski² ; S. Chapeland² ; T. Kleinwort² ; T. Smith² ; V. Bahyl² ; V. Lefebure² ; W. Tomlin²

¹ *Weizman Institute of Science, Israel*

² CERN

³ CERN-IT

This paper describes the evolution of fabric management at CERN's T0/T1 Computing Center, from the selection and adoption of prototypes produced by the European DataGrid (EDG) project[1] to enhancements made to them.

In the last year of the EDG project, developers and service managers have been working to understand and solve operational and scalability issues.

CERN has adopted and strengthened Quattor[2], EDG's installation and configuration management toolsuite, for managing all Linux clusters and servers in the Computing Center, replacing existing legacy management systems. Enhancements to the original prototype include a redundant and scalable server architecture using proxy technology and producing plug-in components for configuring system and LHC computing services.

CERN now coordinates the maintenance of Quattor, making it available to other sites.

Lemon[3], the EDG fabric monitoring framework, has been progressively deployed onto all managed Linux nodes. We have developed sensors to instrument fabric nodes to provide us with complete performance and exception monitoring information.

Performance visualization displays and interfaces to the existing alarm system have also been provided.

LEAF[4], the LHC-Era Automated Fabric toolset, comprises the State Management System, a tool to enable high-level configuration commands to be issued to sets of nodes during both hardware and service management Use Cases, and the Hardware Management System, a tool for administering hardware workflows and for visualizing and locating equipment.

Finally, we will describe issues currently being addressed and planned future developments.

Poster Session 2 - Tableau: 28 / 441

D0 data processing within EDG/LCG

D. Groep¹ ; J. Templon¹ ; K. Bos¹ ; P. Maettig² ; R. Byrom³ ; S. Fisher³ ; T. Harenberg⁴ ; W. Leeuwen¹

¹ NIKHEF

² University of Wuppertal

³ Rutherford Appleton Laboratory

⁴ UNIVERSITY OF WUPPERTAL

Auteur(s) contact: harenberg@physik.uni-wuppertal.de

The D0 experiment at the Tevatron is collecting some 100 Terabytes of data each year and has a very high need of computing resources for the various parts of the physics program. D0 meets these demands by establishing a world - increasingly based on GRID technologies.

Distributed resources are used for D0 MC production and data reprocessing of 1 billion events, requiring 250 TB to be transported over WANs. While in 2003 most of this computing at remote sites was distributed manually, some data reprocessing was performed with the EDG. In 2004 GRID tools are increasingly and successfully employed.

We will report on performing MC production and data reprocessing using EDG and LCG. We will explain how the D0 computing environment was linked to these GRID platforms, and will discuss some lessons

learned (for both Grid computing and preparing applications for distributed operation) from the D0 reprocessing on EDG, subjecting a generic Grid infrastructure to real data for the first time.

An outlook on plans for applying LCG within D0 is given.

Distributed Computing Services / 377

DIRAC - The Distributed MC Production and Analysis for LHCb

A. Pickford¹ ; A. Soroko² ; A. TSAREGORODTSEV³ ; G. Kuznetsov⁴ ; G.N. Patrick⁴ ; I. Korolko⁵ ; I. Stokes-Rees⁶ ; J. Blouw⁷ ; J. Closier⁸ ; J. Saborido-Silva⁹ ; M. Frank⁸ ; M. Gandelman¹⁰ ; M. Sanchez-Garcia⁹ ; M. Tobin¹¹ ; M. Witek¹² ; N. Brook¹³ ; P. Charpentier⁸ ; R. Bernet¹⁴ ; R. Graciani-Diaz¹⁵ ; U. Egede¹⁶ ; V. GARONNE³ ; V. Romanovski¹⁷ ; V. Vagnoni¹⁸

¹ *University of Glasgow*

² *Oxford*

³ *CNRS-IN2P3-CPPM, MARSEILLE*

⁴ *Rutherford Appleton Lab, Chilton*

⁵ *ITEP, Moscow*

⁶ *Particle Physics, Oxford; CPPM, Marseille*

⁷ *GRIDKA, Karlsruhe*

⁸ *CERN*

⁹ *University of Santiago de Compostela*

¹⁰ *University of Rio de Janeiro*

¹¹ *University of Liverpool*

¹² *University of Cracow*

¹³ *H.H. Wills Physics Laboratory, Bristol*

¹⁴ *University of Zurich*

¹⁵ *University of Barcelona*

¹⁶ *Imperial College, London*

¹⁷ *IHEP, Protvino*

¹⁸ *INFN, Bologna*

Auteur(s) contact: atsareg@in2p3.fr

DIRAC is the LHCb distributed computing grid infrastructure for MC production and analysis. Its architecture is based on a set of distributed collaborating services. The service decomposition broadly follows the ARDA project proposal, allowing for the possibility of interchanging the EGEE/ARDA and DIRAC components in the future. Some components developed outside the DIRAC project are already in use as services, for example the File Catalog developed by the AliEn project.

An overview of the DIRAC architecture will be given, in particular the recent developments to support user analysis. The main design choices will be presented. One of the main design goals of DIRAC is the simplicity of installation, configuring and operation of various services. This allows all the DIRAC resources to be easily managed by a single Production Manager.

The modular design of the DIRAC components allows its functionality to be easily extended to include new computing and storage elements or to handle new tasks. The DIRAC system already uses different types of computing resources - from single PC's to a variety of batch system to the Grid

environment. In particular, the use of the LCG2 environment will be presented. Different ways to utilise LCG2 resources will be examined as well as the issue of interoperability between the LCG2 and DIRAC sites.

Poster Session 2 - Tableau: 57 / 368

DIRAC Lightweight information and monitoring services using XML-RPC and Instant Messaging

A. Tsaregorodtsev¹ ; I. Stokes-Rees² ; V. GARONNE³

¹ CPPM, IN2P3

² UNIVERSITY OF OXFORD PARTICLE PHYSICS

³ CNRS-IN2P3-CPPM, MARSEILLE

Auteur(s) contact: ian.stokes-rees@cern.ch

The DIRAC system developed for the CERN LHCb experiment is a grid infrastructure for managing generic simulation and analysis jobs. It enables jobs to be distributed across a variety of computing resources, such as PBS, LSF, BQS, Condor, Globus, LCG, and individual workstations.

A key challenge of distributed service architectures is that there is no single point of control over all components. DIRAC addresses this via two complementary features: a distributed Information System, and an XMPP (Extensible Messaging and Presence Protocol) Instant Messaging framework.

The Information System provides a concept of local and remote information sources.

Any information which is not found locally will be fetched from remote sources. This allows a component to define its own state, while fetching the state of other components directly from those components, or via a central Information Service. We will present the architecture, features, and performance of this system.

XMPP has provided DIRAC with numerous advantages. As an authenticated, robust, lightweight, and scalable asynchronous message passing system, XMPP is used, in addition to XML-RPC, for inter-Service communication, making DIRAC very fault-tolerant, a critical feature when using Service Oriented Architectures. XMPP is also used for monitoring real-time behaviour of the various DIRAC components.

Finally, XMPP provides XML-RPC like facilities which are being developed to provide control channels direct to Services, Agents, and Jobs. We will describe our novel use of Instant Messaging in DIRAC and discuss directions for the future.

Poster Session 2 - Tableau: 25 / 365

DIRAC Workload Management System

A. Tsaregorodtsev¹ ; I. Stokes-Rees² ; V. garonne¹

¹ CPPM-IN2P3 MARSEILLE

² oxford

Auteur(s) contact: garonne@cppm.in2p3.fr

The Workload Management System (WMS) is the core component of the DIRAC distributed MC production and analysis grid of the LHCb experiment. It uses a central Task database which is accessed via a set of central Services with Agents running on each of the LHCb sites. DIRAC uses a 'pull' paradigm where Agents request tasks whenever they detect their local resources are available.

The collaborating central Services allow new components to be plugged in easily. These Services can perform functions such as scheduling optimization, task prioritization, job splitting and merging, to name a few. They provide also job status information for various monitoring clients. We will discuss the services deployment and operation with particular emphasis on the robustness and scalability issues.

The distributed Agents have modular design which allows easy functionality extensions to adapt to the needs of a particular site. The Agent installation have only basic pre-requisites which makes it easy for new sites to be incorporated. An Agent can be deployed on a gatekeeper of a large cluster or just on a single worker node of the LCG grid. PBS,LSF,BQS, Condor,LCG,Globus can be used as the DIRAC computing resources.

The WMS components use XML-RPC and instant messaging Jabber protocols for communication which increases the overall reliability of the system. The jobs handled by the WMS are described using Classad library which facilitates the interoperability with other grids.

Online Computing / 449

DZERO Data Acquisition Monitoring and History Gathering

Auteur(s): G. Watts¹

Co-auteur(s) Andy Haas²

¹ UNIVERSITY OF WASHINGTON

² University of Washington

Auteur(s) contact: gwatts@uw.edu

The DZERO Collider Experiment logs many of its Data Acquisition Monitoring Information in long term storage. This information is most frequently used to understand shift history and efficiency. Approximately two kilobytes of information is stored every 15 second. We describe this system and the web interface provided. The current system is distributed, running on Linux for the back end and Windows for the web interface front end and data logging. We also discuss the development path we have taken for the database backend, from use of root, to Oracle, and back to root.

Distributed Computing Services / 414

Data Management in EGEE

A. Frohner¹ ; G. McCance¹ ; K. Nienartowicz¹ ; P. Badino¹ ; P. Kunszt¹ ; R. Da Rocha¹

¹ CERN

Auteur(s) contact: krzysztof.nienartowicz@cern.ch, gavin.mccance@cern.ch

Data management is one of the cornerstones in the distributed production computing environment that the EGEE project aims to provide for a European e-Science infrastructure. We have designed a set of services based on previous experience in other Grid projects, trying to address the requirements of our user communities.

In this paper we summarize the most fundamental requirements and constraints as well as the security, reliability, stability and robustness considerations that have driven the architecture and the particular choice for service decomposition in our service-oriented architecture. We discuss the interaction of our services with each other, their deployment models and how failures are being managed.

The three service groups for data management services are the Storage Element, the Data Scheduling and the Catalog services. The Storage Element exposes interfaces to Grid managed storage, with the appropriate semantics in the Grid distributed environment. The Catalog services contain all the metadata related to data: The File Catalog maintains a file-system-like view of the files in the Grid in a logical user namespace, the Replica Catalog keeps track of identical copies of the files distributed in different Storage Elements and the Metadata Catalog keeps application specific information about the files. The Data Scheduling services take care of controlled data transfer and keep the information in the Catalog services consistent with what is actually available in the Storage Elements, acting as the binding between the two.

We conclude with first experiences and examples of use-cases for High Energy Physics applications.

Poster Session 2 - Tableau: 68 / 362

Data Rereprocessing on Worldwide Distributed Systems

D. Wicke¹ ; M. Diesburg²

¹ *Fermilab*

² *Fermi National Accelerator Laboratory*

Abstract:

The D0 experiment faces many challenges enabling access to large datasets for physicists on 4 continents. The strategy of solving these problems on worldwide distributed computing clusters is followed.

Already since the begin of TEvatron RunII (March 2001) all Monte-Carlo simulations are produced outside of Fermilab at remote systems. For analyses

as system of regional analysis centers (RACs) was established which supply the associated institutes with the data. This structure which is similar the the Tier structure foreseen for LHC was used in autumn 2003 to rereprocess all D0-data with the uptodate and much improved reconstruction software.

As the first running experiment D0 has implemented and operated all important computing dask of a high energy physics experiment on worldwide distributed systems.

The experiences gained in D0 can be applied to judge the LHC computing model.

Poster Session 2 - Tableau: 33 / 294

Data management services of NorduGrid

A. Konstantinov¹; A. Waananen²; B. Konya³; F. Ould-Saada⁴; J.L. Nielsen²; J.R. Hansen²; M. Ellert⁵; O. Smirnova³; P. Eerola³; T. Ekelof⁶

¹ *Vilnius University, Lithuania*

² *Niels Bohr Institute, Copenhagen, Denmark*

³ *Lund University, Sweden*

⁴ *University of Oslo, Norway*

⁵ *CERN*

⁶ *Uppsala University, Sweden*

In common grid installations, services responsible for storing big data chunks, replication of those data and indexing their availability are usually completely decoupled. And a task of synchronizing data is passed to either user-level tools or separate services (like spiders) which are subject to failure and usually cannot perform properly if one of underlying services fails too.

The NorduGrid Smart Storage Element (SSE) was designed to try to overcome those problems by combining the most desirable features into one service. It uses HTTPS/G for secure data transfer, Web Services for control (through same HTTPS/G channel) and can provide information to indexing services used in middlewares based on the Globus Toolkit (TM). At the moment, those are the Replica Catalog and the Replica Location Service. The modular internal design of the SSE and the power of C++ object programming allows to add support for other indexing services in an easy way.

There are plans to complement it with a Smart Indexing Service capable of resolving inconsistencies hence creating a robust distributed data storage system.

Poster Session 2 - Tableau: 37 / 205

Database Usage and Performance for the Fermilab Run II Experiments

A. Kumar¹; D. Bonham¹; D. Box¹; D. Litvintsev²; E. Gallas¹; E. Wicklund¹; J. Kowalkowski¹; J. Trumbo¹; L. Lueking³; M. Vittone-Wiersma¹; N. Stanfield¹; P. Maksimovic⁴; R. Jetton¹; S. Kovich¹; S.P. White¹; T. Yasuda¹; Y. Guo¹

¹ *Fermilab*

² *FNAL*

³ *FERMILAB*

⁴ *John Hopkins University*

Auteur(s) contact: lueking@fnal.gov

The Run II experiments at Fermilab, CDF and D0, have extensive database needs covering many areas of their online and offline operations. Delivery of the data to users and processing farms based around the world has represented major challenges to both experiments. The range of applications employing databases includes data management, calibration (conditions), trigger information, run configuration, run quality, luminosity, and others. Oracle is the primary database product being used for these applications at Fermilab and some of its advanced features have been employed, such as table partitioning and replication. There is also experience with

open source database products such as MySQL for secondary databases. A general overview of the operation, access patterns, and transaction rates is examined and the potential for growth in the next year presented. The two experiments, while having similar requirements for availability and performance, employ different architectures for database access. Details of the experience for these approaches will be compared and contrasted, as well as the evolution of the delivery systems throughout the run. Tools employed for monitoring the operation and diagnosing problems will also be described.

BOF : Semantic Web applications in HEP / 533

Defining a Semantic Web Initiative for High Energy Physics

B. White¹

¹ *SLAC*

Auteur(s) contact: bebo@slac.stanford.edu

During a recent visit to SLAC, Tim Berners-Lee challenged the High Energy Physics community to identify and implement HEP resources to which Semantic Web technologies could be applied. This challenge comes at a time when a number of other scientific disciplines (for example, bioinformatics and chemistry) have taken a strong initiative in making information resources compatible with Semantic Web technologies and in the development of associated tools and applications.

The CHEP conference series has a strong history of identifying and encouraging adoption of new technologies. The most notable of these technologies include the Web itself and Grid computing. The Semantic Web could have a similar potential.

Topics of discussion in this BoF include (but are not limited to): Definition of the Semantic Web; Semantic Web component technologies; Review of current Semantic Web-related efforts in HEP; Semantic Web resources that are publicly available; What needs to be done next.

Distributed Computing Systems and Experiences / 389

Deploying and operating LHC Computing Grid 2 (LCG2) During Data Challenges

CERN-IT-GD-GIS section¹ ; E. Leonardi² ; M. Schulz¹

¹ *CERN*

² *INFN Rome*

Auteur(s) contact: markus.schulz@cern.ch

LCG2 is a large scale production grid formed by more than 40 worldwide distributed sites. The aggregated number of CPUs exceeds 3000 several MSS systems are integrated in the system. Almost all sites form an independent administrative domain. On most of the larger sites the local computing resources have been integrated into the grid.

The system has been used for large scale production by LHC experiments for several month.

During the operation the software went through several versions and had to be upgraded including non backward compatible upgrades.

We report on the experience gained setting up the service, integrating sites and operating it under the load of the production.

Poster Session 2 - Tableau: 14 / 468

Deployment of SAM for the CDF Experiment

A. Baranovski¹ ; A. Kreymer¹ ; A. Lyon¹ ; A. Sill² ; F. Ratnikov³ ; G. Garzoglio¹ ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter⁴ ; L. Lueking¹ ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber⁷ ; R. Illingworth⁷ ; R. Kennedy¹ ; R. St.Denis⁵ ; S. Belforte⁸ ; S. Stonjek⁹ ; S. Veseli¹ ; S. White¹ ; U. Kerzel¹⁰ ; V. Bartsch⁶ ; W. Merritt¹

¹ *Fermi National Accelerator Laboratory*

² *Texas Tech University*

³ *Rutgers University*

⁴ *Fermilab*

⁵ *Glasgow University*

⁶ *Oxford University*

⁷ *FERMI NATIONAL ACCELERATOR LABORATORY*

⁸ *INFN / Trieste*

⁹ *Fermi National Accelerator Laboratory / University of Oxford*

¹⁰ *Karlsruhe University*

CDF is an experiment at the Tevatron at Fermilab. One dominating factor of the experiments' computing model is the high volume of raw, reconstructed and generated data. The distributed data handling services within SAM move these data to physics analysis applications. The SAM system was already in use at the D-Zero experiment. Due to difference in the computing model of the two experiments some aspects of the SAM system had to be adapted. We will present experiences from the adaptation and the deployment phase. This includes the behavior of the SAM system on batch systems of very different sizes and type as well as the interaction between the datahandling and the storage systems, ranging from disk pools to tape systems. In particular we will cover the problems faced on large scale compute farms. To accommodate the needs of Grid computing, CDF deployed installations consisting of SAM for datahandling and CAF for high throughput batch processing. The CDF experiment already had experiences with the CAF system. We will report on the deployment of the combined system.

Distributed Computing Services / 274

Design and Implementation of a Notification Model for Grid Monitoring Events

G. Rubini¹ ; G. Tortone² ; N. De Bortoli² ; S. Androozzi¹ ; S. Fantinel³

¹ *INFN-CNAF BOLOGNA (ITALY)*

² *INFN - NAPLES (ITALY)*

³ *INFN - LEGNARO (ITALY)*

Auteur(s) contact: nataschia.debortoli@na.infn.it

GridICE is a monitoring service for the Grid, it measures significant Grid related resources parameters in order to analyze usage, behavior and performance of the Grid and/or to detect and notify fault situations, contract violations, and user-defined events. In its first implementation, the notification service relies on a simple model based on a pre-defined set of events.

The growing interest for more flexible and scalable notification capabilities from several LHC experiments has led us to study a more suitable solution satisfying their requirements. With this paper we present both model and design of a notification service which main functionalities are: filtering, transformation, and routing of data. It basically collects a large number of incoming streams of data items from monitored resources (events), filters them according to user profiles or queries describing users information preferences (subscriptions) and finally, after a customization of matched data items, notifies users whose interests are satisfied (event consumers).

Our proposal significantly improves the notification capabilities in current Grid systems by providing flexible means for specifying both topic and content based subscriptions, moreover it provides an efficient matchmaking engine. The new component has been developed and integrated in the GridICE service based on users expressed interests.

Poster Session 1 - Tableau: 22 / 473

Designing a Useful Email System

A. Lilianstrom¹ ; J. Schmidt¹ ; K. Hill¹ ; R. Pasetes¹

¹ *Fermilab*

Email is an essential part of daily work. The FNAL gateways process in excess of 700,000 messages per week. Among those messages are many containing viruses and unwanted spam. This paper outlines the FNAL email system configuration. We will discuss how we have defined our systems to provide optimum uptime as well as protection against viruses, spam and unauthorized users.

Poster Session 3 - Tableau: 63 / 31

Detector-independent vertex reconstruction toolkit (VERTIGO)

G. Richter¹ ; W. Mitaroff¹ ; W. Waltenberger¹

¹ *Austrian Academy of Sciences // Institute of High Energy Physics*

Auteur(s) contact: mitaroff@hephy.oeaw.ac.at

A proposal is made for the design and implementation of a detector-independent vertex reconstruction toolkit and interface to generic objects (VERTIGO). The first stage aims at re-using existing state-of-the-art algorithms for geometric vertex finding and fitting by both linear (Kalman filter) and robust estimation methods. Prototype candidates for the latter are a wide range of adaptive filter algorithms being developed for LHC/CMS, as well as proven ones (like ZVTOP of SLC/SLD). In a second stage, also kinematic constraints will be included for the benefit of complex multi-vertex topologies.

The design is based on modern object-oriented techniques. A core (RAVE) is surrounded by a shell of abstract interfaces (using adaptors for access from/to the particular environment) and a set of analysis and debugging tools. The implementation follows an open source approach and is easily adaptable to future standards.

Work has started with the development of a specialized visualisation tool, following the model-view-controller (MVC) paradigm; it is based on COIN3D and may also include interactivity by PYTHON scripting. A persistency storage solution, intended to provide a general data structure, was originally based on top of ROOT and is currently being extended for AIDA and XML compliance; interfaces to existing or future event reconstruction packages are easily implementable. Flexible linking to a math library is an important requirement; at present we use CLHEP, which could be replaced by a generic product.

Poster Session 3 - Tableau: 12 / 76

Deterministic Annealing for Vertex Finding at CMS

E. Chabanat¹ ; N. Estre¹

¹ IN2P3

Auteur(s) contact: chabanat@ipnl.in2p3.fr

CMS and others LHC experiments offer a new challenge for the vertex reconstruction: the elaboration of efficient algorithms at high-luminosity beam collisions. We present here a new algorithm in the vertex finding field : Deterministic Annealing (DA). This algorithm comes from information theory by analogy to statistical physics and has already been used in clustering and classification problems. In our purpose, the main job is to code information of a set of tracks into prototypes which will be our vertices at the end of the process. The advantages of such a technique is to globally search all vertices at one time and a priori knowledge of the expected number of vertices is not required: the algorithm creates new vertices by a phase transition mechanism which will be describe in this contribution. Thus, the first part of this talk is devoted to a short description of the DA algorithm and to the necessary introduction of the concept of apex points which stand for tracks in this method ; then a discussion of vertex reconstruction efficiencies follows consisting finding DA's internal parameters and making a comparison between DA and the most popular vertex finding algorithm. This comparison is done considering 4000 bbar events generated in the detector central region without pile-up in a first approach ; primary and secondary vertices reconstruction results are shown. Then performances of DA in regional vertex search with regional tracks reconstruction is also presented and lead to a short study of 500 bbar event with pile-up at low luminosity.

Computer Fabrics / 496

Developing & Managing a large Linux farm - the Brookhaven Experience

A. Withers¹ ; C. Hollowell² ; J. Smith¹ ; O. Rind² ; R. Hogue² ; S. Misawa² ; T. Chan¹ ; Tomasz WLODEK³

¹ *BROOKHAVEN NATIONAL LAB*

² *Brookhaven National Lab*

³ *BNL*

Auteur(s) contact: tony@bnl.gov

This presentation describes the experiences and the lessons learned by the RHIC/ATLAS Computing Facility (RACF) in building and managing its 2,700+ CPU (and growing) Linux Farm over the past 6+ years. We describe how hardware cost, end-user needs, infrastructure, footprint, hardware configuration, vendor selection, software support and other considerations have played a role in the process of steering the growth of the RACF Linux Farm, and how they help shape our future hardware purchase decisions. As well as a detailed description of the challenges encountered and of the solutions used in managing and configuring a large, heterogenous Linux Farm (2700+ CPU's) in the midst of an ongoing transition from being a generally local resource to a global, Grid-aware resource within a larger, distributed computing environment is provided.

Distributed Computing Services / 393

Development and use of MonALISA high level monitoring services for the star unified Meta-Scheduler

E. Efstathiadis¹ ; I. Legrand² ; J. Lauret³ ; L Hajdu⁴

¹ *BROOKHAVEN NATIONAL LABORATORY*

² *California Institute of Technology*

³ *Brookhaven National Laboratory*

⁴ *BNL*

Auteur(s) contact: stratos@bnl.gov

As a PPDG cross-team joint project, we proposed to study, develop, implement and evaluate a set of tools that allow Meta-Schedulers to take advantage of consistent information (such as information needed for complex decision making mechanisms) across both local and/or Grid Resource Management Systems (RMS).

We will present and define the requirements and schema by which one can consistently provide queue attributes for the most common batch systems (PBS, LSF, Condor, SGE, etc). We evaluate the best scalable and lightweight approach to access the monitored parameters from a client perspective and, in particular, the feasibility of accessing real-time and aggregate information using the MonALISA monitoring framework. Client programs are envisioned to function in a non-centralized, fault tolerant fashion. Inherent delays as well as scalability issues of each approach (implementing it at a large number of sites) will be discussed.

The MonALISA monitoring framework, being an ensemble of autonomous multi-threaded, agent based systems which are registered as dynamic services and are able to collaborate and cooperate in performing a wide range of monitoring tasks in a large scale distributed applications, is a natural choice for such a project. MonALISA is designed to easily integrate existing monitoring tools and procedures and provide information in a dynamic self-describing way to any other service or client. We intend to demonstrate the usefulness of this consistent approach for queue monitoring by implementing a monitoring agent within the STAR Unified Meta-Scheduler (SUMS) framework.

We believe that such developments could highly benefit Grid laboratory efforts such as the Grid3+ and the OpenScience Grid (OSG).

Poster Session 3 - Tableau: 13 / 32

Development of algorithms for cluster finding and track reconstruction in the forward muon spectrometer of ALICE experiment

A. Vodopianov¹ ; A. Zinchenko² ; G. Chabratova¹ ; V. Pismennaia¹

¹ *Joint Institute for Nuclear Research, Dubna, Russia*

² *JOINT INSTITUTE FOR NUCLEAR RESEARCH (JINR), DUBNA, RUSSIA*

Auteur(s) contact: alexandre.zintchenko@cern.ch

A simultaneous track finding / fitting procedure based on Kalman filtering approach has been developed for the forward muon spectrometer of ALICE experiment.

In order to improve the performance of the method in high-background conditions of the heavy ion collisions the “canonical” Kalman filter has been modified and supplemented by a “smoother” part. It is shown that the resulting “extended” Kalman filter gives better tracking results and offers higher flexibility.

To further improve the tracking performance in a high occupancy environment a new algorithm for cluster / hit finding in cathode pad chambers of the muon spectrometer has been developed. It is based on the expectation maximization procedure for a shape deconvolution of overlapped clusters.

It is demonstrated that the proposed method allows to reduce the loss of the coordinate reconstruction accuracy for high hit multiplicities and achieve better tracking results.

Both the hit finding and track reconstruction algorithms have been implemented within the AliRoot software framework.

Core Software / 268

Developments of Mathematical Software Libraries for the LHC experiments

Auteur(s): L. Moneta¹

Co-auteur(s) A Zsenei¹ ; F. James¹ ; M. Hatlo¹ ; M. Winkler¹ ; P. Mato Vila¹

¹ *CERN*

Auteur(s) contact: lorenzo.moneta@cern.ch

The main objective of the MathLib project is to give expertise and support to the LHC experiments on mathematical and statistical computational methods. The aim is to provide a coherent set of mathematical libraries. Users of this set of libraries are developers of experiment reconstruction and simulation software, of analysis tools frameworks, such as ROOT, and physicists performing data analysis.

After having performed a detailed evaluation of the existing functionality present in GSL, a general purpose mathematical library, and in more HEP specific libraries such as CLHEP, CERNLIB and ROOT, a new object oriented library has been started to be developed. The new library incorporates or uses most of the functions and algorithms of the already existing libraries. Examples of these functions and algorithms are mathematical special functions, linear algebra, minimization and any other required numerical algorithms. Wrappers to these are written in C++ and integrated in a coherent object oriented framework. Interfaces to the Python interactive environment are as well provided.

An overview of the project activities will be presented, describing in detail the current functionality of the library and its design. Furthermore, the object oriented implementation of Minuit, a fitting and minimization framework, will be covered in the presentation.

Computer Fabrics / 325

Disk storage technology for the LHC T0/T1 centre at CERN

A. Horvath¹ ; H. Meinhard¹ ; P. Kelemen¹

¹ CERN-IT

Auteur(s) contact: helge.meinhard@cern.ch

By 2008, the T0/T1 centre for the LHC at CERN is estimated to use about 5000 TB of disk storage. This is a very significant increase over the about 250 TB running now. In order to be affordable, the chosen technology must provide the required performance and at the same time be cost-effective and easy to operate and use.

We will present an analysis of the cost (both in terms of material and personnel) of the current implementation (network-attached storage), and then describe detailed performance studies with hardware currently in use at CERN in different configurations of filesystems on software or hardware RAID arrays over disks. Alternative technologies that have been evaluated by CERN in varying depth (such as arrays of SATA disks with a Fiber Channel uplink, distributed disk storage across worker nodes, iSCSI solutions, SANFS, ...) will be discussed. We will conclude with an outlook of the next steps to be taken at CERN towards defining the future disk storage model.

Distributed Computing Systems and Experiences / 497

Distributed Computing Grid Experiences in CMS DC04

A. Anzar¹ ; A. Fanfani² ; A. Heavey¹ ; A. Sciaba³ ; A. Zahn⁴ ; B. MacEvoy⁵ ; B. Martelli⁶ ; C. Grandi⁷ ; C. Prescott⁴ ; D. Bonacorsi⁶ ; D. Bradley⁸ ; D. Colling⁹ ; F. Donno³ ; F. Fanzago¹⁰ ; G. Donvito¹¹ ; G. Graham¹ ; G. Maggi¹¹ ; H. Tallini⁹ ; I. Fisk¹ ; I. Legrand¹² ; J. Hernandez¹³ ; J. Kaiser¹ ; J. Letts¹⁴ ; J. Rodriguez⁴ ; J. Weigand¹ ; K. Bockjoo⁴ ; L

Wakefield⁹ ; L. Dell'Agnello⁶ ; L. Giacchetti¹ ; L. Tuura¹⁵ ; M. Biasotto¹⁶ ; M. Corvo¹⁰ ; M. Ernst¹⁷ ; M. Mazzucato¹⁰ ; N. De Filippis¹¹ ; N. Kuropatine¹ ; N. Ratnikova¹ ; N. Sinanis¹⁸ ; P. Capiluppi⁷ ; P. GARCIA-ABIA¹⁹ ; R. Pordes¹ ; S. Campana³ ; S. Fantinel¹⁶ ; S. Suresh¹² ; T. Martin¹⁴ ; T. Perelmutov¹ ; W. Jank¹⁸ ; Y. Wu¹

¹ FERMILAB

² INFN-BOLOGNA (ITALY)

³ EIS CERN

⁴ UFL

⁵ IMPERIAL COLLEGE LONDON

⁶ INFN-CNAF

⁷ INFN-Bologna

⁸ University of Wisconsin

⁹ Imperial College London

¹⁰ INFN-Padova

¹¹ INFN-Bari

¹² Caltech

¹³ CIEMAT Madrid

¹⁴ UCSD

¹⁵ Northeastern University, Boston

¹⁶ INFN-Legnaro

¹⁷ DESY

¹⁸ CERN

¹⁹ CIEMAT

Auteur(s) contact: alessandra.fanfani@cern.ch, ian.fisk@cern.ch

In March-April 2004 the CMS experiment undertook a Data Challenge(DC04). During the previous 8 months CMS undertook a large simulated event production. The goal of the challenge was to run CMS reconstruction for sustained period at 25Hz input rate, distribute the data to the CMS Tier-1 centers and analyze them at remote sites. Grid environments developed in Europe by the LHC Computing Grid (LCG) in Europe and in the US with Grid2003 were utilized to complete the aspects of the challenge.

During the simulation phase, US-CMS utilized Grid2003 to simulate and process approximately 17 million events. Simultaneous usage of CPU resources peaked at 1200 CPUs, controlled by a single FTE. Using Grid3 was a milestone for CMS computing in reaching a new magnitude in the number of autonomously cooperating computing sites for production. The use of Grid-based job execution resulted in reducing the overall support effort required to submit and monitor jobs by a factor of two.

During the challenge itself, the CMS groups from Italy and Spain used the LCG Grid Environment to satisfy challenge requirements. The LCG Replica Manager was used to transfer the data. The CERN RLS provided the needed replica catalogue functionality. The LCG submission system based on the Resource Broker was used to submit analysis jobs to the sites hosting the data. A CMS dedicated GridICE monitoring was activated to monitor both services and resources.

A description of the experiences, successes and lessons learned from both experiences with grid infrastructure is presented.

Distributed Filesystem Evaluation and Deployment at the US-CMS Tier-1 Center

D. Fagan¹ ; J. kaiser¹ ; L. Lisa Giacchetti¹ ; M. Ernst² ; R. Pasetes¹

¹ *FERMILAB*

² *DESY*

Auteur(s) contact: michael.ernst@desy.de

The scalable serving of shared filesystems across large clusters of computing resources continues to be a difficult problem in high energy physics computing. The US CMS group at Fermilab has performed a detailed evaluation of hardware and software solutions to allow filesystem access to data from computing systems.

The goal of the evaluation was to arrive at a solution that was able to meet the growing needs of the US-CMS Tier-1 facility. The system needed to be scalable and be able to grow with the increasing size of the facility, load balanced and with high performance for data access, reliable and redundant with protection against failures, and manageable and supportable given a reasonable level of effort.

Over the course of a one year evaluation the group developed a suite of tools to analysis performance and reliability under load conditions, and then applied these tools to evaluations systems at Fermilab. In this presentation we will describe the suite of tools developed, the results of the evaluation process, the system and architecture that were eventually chosen, and the experience so far supporting a user community.

Poster Session 2 - Tableau: 27 / 276

Distributed Testing Infrastructure and Processes for the EGEE Grid Middleware

D. Bosio¹ ; D. Collados¹ ; D. Groep² ; D. Salomoni² ; J. Templon² ; L. Guy¹ ; M. Reale¹ ; M. Theile¹ ; S. Traylen³

¹ *CERN*

² *NIKHEF*

³ *RAL*

Auteur(s) contact: leanne.guy@cern.ch

Extensive and thorough testing of the EGEE middleware is essential to ensure that a production quality Grid can be deployed on a large scale as well as across the broad range of heterogeneous resources that make up the hundreds of Grid computing centres both in Europe and worldwide.

Testing of the EGEE middleware encompasses the tasks of both verification and validation. In addition we test the integrated middleware for stability, platform independence, stress resilience, scalability and performance.

The EGEE testing infrastructure is distributed across three major EGEE grid centres in three countries: CERN, NIKHEF and RAL.

As much as is possible the testing procedures are automated and integrated with the EGEE build system. This allows for continuous testing together with the incremental daily code builds, fast and early feedback to developers of bug, and for the easy inclusion of regression tests.

This paper will report on the initial results of the testing procedures, frameworks and automation techniques adopted by the EGEE project, the advantages and disadvantages of test automation and the issues involved in testing a complex distributed middleware system in a distributed environment.

Poster Session 2 - Tableau: 40 / 419

Distributed Tracking, Storage, and Re-use of Job State Information on the Grid

A. GUARISE¹; A. Gianelle²; A. Krenek³; A. MARASCHINI⁴; A. Terracina⁵; D. Kouril³; D. Rebatto⁶; E. Ronchieri⁷; F. GIACOMINI⁸; F. Pacini⁵; F. Prelz⁶; G. Avellino⁵; G. Patania⁹; J. Pospisil³; J. Sitera³; J. Skrabal³; L. Matyska³; L. Zangrando²; M. MARCHI¹⁰; M. MEZZADRI¹⁰; M. MORDACCHINI¹¹; M. MULAC¹²; M. Pappalardo¹³; M. Ruda³; M. Sgaravatto²; M. Vocu³; P. Andreetto²; S. Andreozzi⁷; S. Beco⁵; S. Borgia²; S. MONFORTE¹⁴; V. Ciaschini⁷; Z. Salvat³

¹ INFN Torino

² INFN Padova, Italy

³ CESNET, CZECH REPUBLIC

⁴ DATAMAT

⁵ DATAMAT, Italy

⁶ INFN Milano, Italy

⁷ INFN Cnaf, Italy

⁸ INFN Cnaf

⁹ INFN Torino, Italy

¹⁰ INFN Milano

¹¹ INFN Padova

¹² CESNET

¹³ INFN Catania, Italy

¹⁴ INFN Catania

Auteur(s) contact: ljocha@ics.muni.cz

The Logging and Bookkeeping service tracks job passing through the Grid. It collects important events generated by both the grid middleware components and applications, and processes them at a chosen L&B server to provide the job state. The events are transported through secure reliable channels. Job tracking is fully distributed and does not depend on a single information source, the robustness is achieved through speculative job state computation in case of reordered, delayed or lost events. The state computation is easily adaptable to modified job control flow.

The events are also passed to the related Job Provenance service. Its purpose is a long-term storage of information on job execution, environment, and the executable and input sandbox files. The data can be used for debugging, post-mortem analysis, or re-running jobs. The data are kept by the job-provenance storage service in a compressed format, accessible on per-job basis. A complementary index service is able to find particular jobs

according to configurable criteria, e.g. submission time or “tags” assigned by the user. A user client to support job re-execution is planned.

Both the L&B and Job Provenance index server provide web-service interfaces for querying. Those interfaces comply with the On-demand producer specification of the R-GMA infrastructure. Hence R-GMA capabilities can be utilized to perform complex distributed queries across multiple servers. Also, aggregate information about job collections can be easily provided.

The L&B service was deployed in the EU DataGrid and Cern LCG projects, the Job Provenance will be deployed in the EGEE project.

Poster Session 2 - Tableau: 46 / 349

Distributed computing and oncological radiotherapy: technology transfer from HEP and experience with prototype systems

F. Foppiano¹ ; J. Moscicki² ; M.G. Pia³ ; S. Guatelli⁴

¹ *IST Genova*

² *CERN*

³ *INFN GENOVA*

⁴ *INFN Genova, Italy*

Auteur(s) contact: maria.grazia.pia@cern.ch

We show how nowadays it is possible to achieve the goal of accuracy and fast computation response in radiotherapeutic dosimetry using Monte Carlo methods, together with a distributed computing model.

Monte Carlo methods have never been used in clinical practice because, even if they are more accurate than available commercial software, the calculation time needed to accumulate sufficient statistics is too long for a realistic use in radiotherapeutic treatment.

We present a complete, fully functional prototype dosimetric system for radiotherapy, integrating various components based on HEP software systems: a Geant4-based simulation, an AIDA-based dosimetric analysis, a web-based user interface, and distributed processing either on a local computing farm or on geographically spread nodes.

The performance of the dosimetric system has been studied in three execution modes: sequential on a single dedicated machine, parallel on a dedicated computing farm, parallel on a grid test-bed. An intermediate software layer, the DIANE system, makes the three execution modes completely transparent to the user, allowing to use the same code in any of the three configurations. Thanks to the integration in a grid environment, any hospital, even small ones or in less wealthy countries, that could not afford the high costs of commercial treatment planning software, may get the chance of using advanced software tools for oncological therapy, by accessing distributed computing resources, shared with other hospitals and institutes belonging to the same virtual organization

Poster Session 3 - Tableau: 35 / 94

Domain Specific Visual Query Language for HEP analysis or How far can we go with user friendliness?

Auteur(s): V M. Moreira do Amaral¹

Co-auteur(s) G. Moerkotte¹ ; S. Helmer¹

¹ *UNIVERSITY OF MANNHEIM*

Auteur(s) contact: amaral@pi3.informatik.uni-mannheim.de

There is a permanent quest for user friendliness in HEP Analysis. This growing need is directly proportional to the analysis frameworks' interface complexity. In fact, the user is provided with an analysis framework that makes use of a General Purpose Language to program the query algorithms. Usually the user finds this overwhelming, since he or she is presented with the complexity of the intricacies of the systems. This way the final user of HEP experiments becomes a forced programmer or an application developer.

In our opinion this inflicts directly or indirectly in the query system performances. For this reason we have decided to invest in a line of research to find a solution that balances the complexity and variability of the analysis queries with the need for simpler query systems interfaces. The ultimate goal is to save time on query algorithms production and to have a way to increase efficiency.

In this communication we are going to present how we explored the hypothesis of generating a visual query language specific for the HEP high level analysis domain. The prototyped framework developed so far, PHEASANT, is supporting our arguments in the feasibility of this approach. Therefore, like in any young Human Centric development project, this raises the need of a broad discussion in order to validate it. We believe to be opening an new fruitful research topic among the community and we expect to motivate both computer science and physics experts into the same discussion.

Distributed Computing Services / 142

Don Quijote - Data Management for the ATLAS Automatic Production System

M. Branco¹

¹ *CERN*

Auteur(s) contact: miguel.branco@cern.ch

As part of the ATLAS Data Challenges 2 (DC2), an automatic production system was introduced and with it a new data management component.

The data management tools used for previous Data Challenges were built as separate components from the existing Grid middleware. These tools relied on a database of its own which acted as a replica catalog.

With the extensive use of Grid technology expected for the most part of the DC2 production, no longer can a data management tool be independent of the Grid middleware. Each Grid relies on its own replica catalog and not on an ATLAS specific tool.

ATLAS DC will attempt to use uniformly the resources provided by three Grids: NorduGrid, US Grid3 and LCG-2. Lecagy system will be supported as well.

The proposed solution was to build a data management proxy system which consists of a common high-level interface, whose implementation depends on each Grid's replica and

metadata catalog as well as the storage backend (mainly “classic” GridFTP servers and SRM).

Don Quijote provides management of replicas in a services oriented architecture, across the several “flavours” of Grid middleware used by ATLAS DC.

With a higher-level interface common across several Grids (and legacy systems) a user (such as the new automatic production system) can seamlessly manage replicas independently of their hosting environment. Given the services-based architecture, a lightweight command line tool is capable of interacting uniformly within each Grid and between Grids (e.g. moving files from LCG-2 to US Grid 3 while maintaining attributes such as the Global Unique Identifier).

Poster Session 3 - Tableau: 40 / 310

Dynamic Matched Filters for Gravitational Waves Detection

A. ELEUTERI¹ ; F. ACERNESE¹ ; F. BARONE² ; G. RUSSO³ ; L. MILANO⁴ ; R. DE ROSA¹ ; S. Pardi⁵

¹ *Napoli University and INFN Napoli, Italy*

² *INFN Napoli and Salerno University, Italy*

³ *Napoli University, Italy*

⁴ *Dipartimento di Scienze Fisiche - Università di Napoli Federico II*

⁵ *DIPARTIMENTO DI MATEMATICA ED APPLICAZIONI "R.CACCIOPPOLI"*

Auteur(s) contact: spardi@na.infn.it

The algorithms for the detection of gravitational waves are usually very complex due to the low signal to noise ratio. In particular the search for signals coming from coalescing binary systems can be very demanding in terms of computing power, like in the case of the classical Standard Matched Filter Technique. To overcome this problem, we tested a Dynamic Matched Filter Technique, still based on Matched Filters, whose main advantage is the requirement of a lower computing power. In this work this technique is described, together with its possible application as a pre-data analysis algorithm. Also the results on simulated data are reported.

Plenary / 535

EU Grid Research - Projects and Vision

M. Lemke^{None}

The European Grid Research vision as set out in the Information Society Technologies Work Programmes of the EU's Sixth Research Framework Programme is to advance, consolidate and mature Grid technologies for widespread e-science, industrial, business and societal use. A batch of Grid research projects with 52 Million EUR EU support was launched during the European Grid Technology Days 15 - 17 September 2004. The portfolio of projects has the potential for turning Europe's strong competence and critical mass in Grid Research into competitive advantages. In this presentation, the Grid research vision of the programme and the new project portfolio will be introduced. More information: www.cordis.lu/ist/grids.

Plenary / 485

Enterasys - Networks that Know

J. ROESE^{None}

Today and in the future businesses need an intelligent network. And Enterasys has the smarter solution. Our active network uses a combination of context-based and embedded security technologies - as well as the industry's first automated response capability - so it can manage who is using your network. Our solution also protects the entire enterprise - from the edge, through the distribution layer, and into the core of the network. Threats are recognized and isolated at the user level, rather than taking your entire network down. It even has the ability to coexist with and enhance your legacy data networking infrastructure and existing security appliances - regardless of the vendor. By continually offering a context-based analysis of network traffic, our solution allows you to see not only what the problem is, but also where it is, and who caused it. And, with the industry's most advanced controls, we're the first solution that's able to resolve threats across the entire network - dynamically or on demand.

Grid Security / 224

Evaluation of Grid Security Solutions using Common Criteria

M. RIGUIDEL¹ ; S. NAQVI¹

¹ TELECOM PARIS

Auteur(s) contact: naqvi@enst.fr

In the evolution of computational grids, security threats were overlooked in the desire to implement a high performance distributed computational system. But now the growing size and profile of the grid require comprehensive security solutions as they are critical to the success of the endeavour. A comprehensive security system, capable of responding to any attack on grid resources, is indispensable to guarantee its anticipated adoption by both the users and the resource providers. Some security teams have started working on establishing in-depth security solutions. The evaluation of their grid security solutions requires excellent criteria to assure sufficient security to meet the needs of its users and resource providers. Grid community's lack of experience in the exercise of the Common Criteria (CC), which was adopted in 1999 as an international standard for security product evaluation, makes it imperative that efforts be exerted to investigate the prospective influence of the CC in advancing the state of grid security. This article highlights the contribution of the CC to establishing confidence in grid security, which is still in need of considerable attention from its designers. The process of security evaluation is outlined and the roles each part of the evaluation may play in obtaining confidence are examined.

Event Processing / 105

Event Data Model in ATLAS

E. MOYSE¹ ; F. Akesson²

¹ *CERN, Geneva, Switzerland*

² *CERN*

Auteur(s) contact: fredrik.akesson@cern.ch

The event data model (EDM) of the ATLAS experiment is presented. For large collaborations like the ATLAS experiment common interfaces and data objects are a necessity to insure easy maintenance and coherence of the experiments software platform over a long period of time. The ATLAS EDM improves commonality across the detector subsystems and subgroups such as trigger, test beam reconstruction, combined event reconstruction, and physics analysis. The object oriented approach in the description of the detector data allows the possibility to have one common raw data flow. Furthermore the EDM allows the use of common software between online data processing and offline reconstruction. One important component of the ATLAS EDM is a common track class which is used for combined track reconstruction across the innermost tracking subdetectors and is also used for tracking in the muon detectors. The structure of the track object and the variety of track parameters are presented. For the combined event reconstruction a common particle class is introduced which serves as the interface between event reconstruction and physics analysis.

Core Software / 199

EventStore: Managing Event Versioning and Data Partitioning using Legacy Data Formats

Auteur(s): C. Jones¹

Co-auteur(s) D. Riley¹ ; G. Sharp¹ ; V. Kuznetsov¹

¹ *CORNELL UNIVERSITY*

HEP analysis is an iterative process. It is critical that in each iteration the physicist's analysis job accesses the same information as previous iterations (unless explicitly told to do otherwise). This becomes problematic after the data has been reconstructed several times. In addition, when starting a new analysis, physicists normally want to use the most recent version of reconstruction. Such version control is useful for data managed by a single physicist using a laptop or small groups of physicists at a remote institution in addition to the collaboration wide managed data.

In this presentation we will discuss our implementation of the EventStore which uses a data location, indexing and versioning service to manage legacy data formats (e.g. an experiment's existing proprietary file format or Root files). A plug-in architecture is used to support adding additional file formats. The core of the system is used to implement three different sizes of services: personal, group and collaboration.

Plenary / 527

Evolution and Revolution in the Design of Computers Based on Nanoelectronics

S. WILLIAMS^{None}

Today's computers are roughly a factor of one billion less efficient at doing their job than the laws of fundamental physics state that they could be. How much of this efficiency gain will we actually be able to harvest? What are the biggest obstacles to achieving many orders of magnitude improvement in our computing hardware, rather than the roughly factor of two we are used to seeing with each new generation of chip? Shrinking components to the nanoscale offers both potential advantages and severe challenges. The transition from classical mechanics to quantum mechanics is a major issue. Others are the problems of defect and fault tolerance: defects are manufacturing mistakes or components that irreversibly break over time and faults are transient interruptions that occur during operation. Both of these issues become bigger problems as component sizes shrink and the number of components scales up massively. In 1955, John von Neumann showed that a completely general approach to building a reliable machine from unreliable components would require a redundancy overhead of at least 10,000 - this would completely negate any advantages of building at the nanoscale. We have been examining a variety of defect and fault tolerant techniques that are specific to particular structures or functions, and are vastly more efficient for their particular task than the general approach of von Neumann. Our strategy is to layer these techniques on top of each other to achieve high system reliability even with component reliability of no more than 97% or so, and a total redundancy of less than 3. This strategy preserves the advantages of nanoscale electronics with a relatively modest overhead.

Distributed Computing Services / 278

Evolution of LCG-2 Data Management

J-P. Baud¹ ; J. Casey¹

¹ CERN

Auteur(s) contact: james.casey@cern.ch

LCG-2 is the collective name for the set of middleware released for use on the LHC Computing Grid in December 2003. This middleware, based on LCG-1, had already several improvements in the Data Management area. These included the introduction of the Grid File Access Library(GFAL), a POSIX-like I/O Interface, along with MSS integration via the Storage Resource Manager(SRM)interface.

LCG-2 was used in the Spring 2004 data challenges by all four LHC experiments. This produced the first useful feedback on scalability and functionality problems in the middleware, especially with regards to data management.

One of the key goals for the Data Challenges in 2004 is to show that the LCG can handle the data for the LHC, even if the computing model is still quite simple. In light of the feedback from the data challenges, and in conjunction with the LHC experiments, a strategy for the improvements required in the data management area was developed. The aim of these improvements was to allow both easier interaction and better performance from the experiment frameworks and other middleware such as POOL.

In this talk, we will first introduce the design of the current data management solution in LCG-2. We will cover the problems and issues

highlighted by the data challenges, as well as the strategy for the required improvements to allow LCG-2 to handle effectively data management at LCG volumes. In particular, we will highlight the new APIs provided, and the integration of GFAL and the EDG Replica Manager functionality with ROOT.

Poster Session 2 - Tableau: 41 / 114

Experience integrating a General Information System API in LCG Job Management and Monitoring Services

A. Delgado Peris¹ ; A. Sciaba¹ ; F. Donno¹ ; P. Mendez Lorenzo¹ ; R. Santinelli¹ ; S. Campana¹

¹ CERN IT/GD

Auteur(s) contact: patricia.mendez@cern.ch

In a Grid environment, the access to information on system resources is a necessity in order to perform common tasks such as matching job requirements with available resources, accessing files or presenting monitoring information. Thus both middleware service, like workload and data management, and applications, like monitoring tools, require an interface to the Grid information service which provides that data.

Even though a unique schema for the published information is defined, actual implementations use different data models, and define different access protocols. Applications interacting with the information service must therefore deal with several APIs, and be aware of the underlying technology in order to use the appropriate syntax for their queries or to publish new information.

We have produced a new high level C++ API that accommodates several existing implementations of the information service such as Globus MDS(LDAP based), MDS3(XML based) and R-GMA(SQL based). It allows applications to access information in a transparent manner loading the needed implementation specific library on demand. Features allowing for the adding and removal of dynamic information have been included as well. A general query language to make the API compatible with future protocols has been used.

In this paper we described the design of this API and the results obtained integrating this API in the Workload Management system and in the GridIce monitoring system of LCG.

Distributed Computing Systems and Experiences / 38

Experience producing simulated events for the DZero experiment on the SAM-Grid

A. Nishandar¹ ; G. Garzoglio² ; I. Terekhov² ; J. Snow³ ; S. Jain¹

¹ University of Texas at Arlington

² FERMI NATIONAL ACCELERATOR LABORATORY

³ Langston University

Auteur(s) contact: kennedy@fnal.gov

Most of the simulated events for the DZero experiment at Fermilab have been historically produced by the “remote” collaborating institutions. One of the principal challenges reported concerns the maintenance of the local software infrastructure, which is generally different from site to site. As the understanding

of the community on distributed computing over distributively owned and shared resources progresses, it becomes increasingly interesting the adoption of grid technologies to address the production of monte-carlo events for high energy physics experiments. The SAM-Grid is a software system developed at Fermilab, which integrates standard grid technologies for job and information management with SAM, the data handling system of the DZero and CDF experiments. During the past few months, this grid system has been tailored for the monte-carlo production of DZero. Since the initial phase of deployment, this experience has exposed an interesting series of requirements to the SAM-Grid services, the standard middleware, the resources and their management and to the analysis framework of the experiment. As of today, the inefficiency due to the grid infrastructure has been reduced to as little as 1%. In this paper, we present our statistics and the “lesson learned” in running large high energy physics applications on a grid infrastructure.

Poster Session 1 - Tableau: 15 / 166

Experience with CORBA communication middleware in the ATLAS DAQ

A. AMORIM¹; A. KAZAROV²; D. BURCKHART-CHROMEK³; D. KLOSE¹; D. Liko⁴; E. BADESCU⁵; I. ALEXANDROV⁶; I. SOLOVIEV⁷; J. FLAMMER³; L. MAPELLI³; L. PEDRO¹; M. CAPRINI⁵; M. DOBSON³; M. Mineev⁸; N. FIUZA DE BARROS¹; R. JONES³; S. KOROBV⁶; S. Kolos⁴; V. Kotov⁸; Y. RYABOV⁷

¹ *Universidade de Lisboa, Portugal*

² *Petersburg Nuclear Physics Institute (PNPI)*

³ *CERN, Geneva, Switzerland*

⁴ *CERN*

⁵ *National Institute of Physics and Nuclear Engineering, Bucharest, Romania*

⁶ *JINR, Dubna, Russia*

⁷ *Petersburg Nuclear Physics Institute, Gatchina, Russia*

⁸ *Joint Institute for Nuclear Research (JINR)*

Auteur(s) contact: sergei.kolos@cern.ch

As modern High Energy Physics (HEP) experiments require more distributed computing power to fulfill their demands, the need for an efficient distributed online services for control, configuration and monitoring in such experiments becomes increasingly important. This paper describes the experience of using standard Common Object Request Broker Architecture (CORBA) middleware for providing a high performance and scalable software, which will be used for the online control, configuration and monitoring in the ATLAS Data Acquisition (DAQ) system. It also presents the experience, which was gained from using several CORBA implementations and replacing one CORBA broker with another.

Finally the paper introduces results of the large scale tests, which have been done on the cluster of more than 300 nodes, demonstrating the performance and scalability of the ATLAS DAQ online services. These results show that the CORBA standard is truly appropriate for the highly efficient online distributed computing in the area of modern HEP experiments.

Poster Session 2 - Tableau: 69 / 185

Experience with Deployment and Operation of the ATLAS Production System and the Grid3+ Infrastructure at Brookhaven National Lab

A. Chan¹ ; D. Yu¹ ; G. Carcassi¹ ; J. Smith¹ ; R. Baker¹ ; T. Wlodek¹ ; W. Deng¹ ; X. Zhao¹ ; Z. Liu¹

¹ *Brookhaven National Laboratory*

Auteur(s) contact: xzhao@bnl.gov

This paper describes the deployment and configuration of the production system for ATLAS Data Challenge 2 starting in May 2004, at Brookhaven National Laboratory, which is the Tier1 center in the United States for the International ATLAS experiment. We will discuss the installation of Windmill (supervisor) and Capone (executor) software packages on the submission host and the relevant security issues. The Grid3+ infrastructure and information service are used for the deployment of grid enabled ATLAS transformations on the Grid3+ computing elements. The Tier 1 hardware configuration includes 95 dual processor Linux compute nodes, 24 TB of NFS disk and an HPSS mass storage system. VOMS server maintains both VO services for US ATLAS and BNL local site policies. This paper describes the work of optimizing the performance and efficiency of this configuration.

Distributed Computing Services / 383

Experience with POOL from the LCG Data Challenges of the three LHC experiments

A. Vaniachine¹ ; C. Cioffi² ; C. Grandi³ ; D. Duellmann⁴ ; D. Malon¹ ; G. Govi⁴ ; H. Schmuecker⁴ ; I. Papadopoulos⁴ ; J. Moscicki⁴ ; K. Karr¹ ; L. GOOSSENS⁴ ; L. Tuura⁵ ; M. Branco⁴ ; M. Frank⁴ ; M. Girone⁴ ; R. Chytraccek⁴ ; T. Barrass⁶ ; V. Innocente⁷ ; W. Tannenbaum⁸ ; Z. Xie⁹

¹ *Argonne National Lab*

² *University of Oxford*

³ *INFN-Bologna*

⁴ *CERN*

⁵ *Northeastern University*

⁶ *University of Bristol*

⁷ *CERN, PH/SFT*

⁸ *Fermilab*

⁹ *Princeton University*

Auteur(s) contact: dirk.duellmann@cern.ch

This presentation will summarise the deployment experience gained with POOL during the first larger LHC experiments data challenges performed. In particular we discuss the storage access performance and optimisations, the integration issues with grid middleware services such as the LCG Replica Location Service (RLS) and the LCG Replica Manager and experience with the POOL proposed way of exchanging meta data (such as File Catalog catalogue entries) in a de-coupled production system.

Online Computing / 209

Experience with Real Time Event Reconstruction Farm for Belle Experiment

I. Adachi¹ ; M. Nakao¹ ; N. Katayama¹ ; R. Itoh¹ ; S. Suzuki¹ ; T. Higuchi¹

¹ KEK

Auteur(s) contact: ryosuke.itoh@kek.jp

A sizeable increase in the machine luminosity of KEKB accelerator is expected in coming years. This may result in a shortage in the data storage resource for the Belle experiment in the near future and it is desired to reduce the data flow as much as possible before writing the data to the storage device.

For this purpose, a realtime event reconstruction farm has been installed in the Belle DAQ system. The farm consists of 60 linux-operated PC servers with dual CPUs. Every event from the event builder is distributed to one of the servers through a socket connection. A full event reconstruction is done on each server so that a sophisticated event selection can be performed to reduce the data flow. The same event reconstruction program as that used in the offline DST production runs on each farm server. Selected events are collected through socket connections and written to a fast disk array.

The farm has been being operated in the beam runs from the beginning of this year and processing the data at an average L1 trigger rate of 450Hz. The experience of the operation is reported at the conference. In particular, the performance of the full event reconstruction and selection is discussed in detail. A scheme to monitor the quality of processed data in real time is also described.

Poster Session 3 - Tableau: 45 / 352

Experience with the Unified Process

M.G. Pia¹

¹ INFN GENOVA

Auteur(s) contact: maria.grazia.pia@cern.ch

The adoption of a rigorous software process is well known to represent a key factor for the quality of the software product and the most effective usage of the human resources available to a software project.

The Unified Process, in particular its commercial packaging known as the RUP (Rational Unified Process) has been one of the most widely used software process models in the software industry for a number of years.

We present the application of the Unified Process and of the RUP to a variety of software projects in the High Energy Physics environment. We illustrate how the UP/RUP provide a flexible process framework, that can be tailored to the different needs of individual software projects. We describe the experience with different approaches (top-down and bottom-up) to the implementation of the process in software organizations.

We document a critical analysis of the effects of the adoption of the UP/RUP, and discuss the relative benefits of the public (UP) and commercial (RUP) versions of the process.

Finally, we discuss the curious results of the effects of applying the RUP to a software development environment that is not aware of adopting it

Poster Session 1 - Tableau: 8 / 474

Experiences Building a Distributed Monitoring System

J. Fromm¹ ; K. Genser¹ ; M. Mengel¹ ; T. Levshina² ; V. Podstavkov¹

¹ *Fermilab*

² *FERMILAB*

The NGOP Monitoring Project at FNAL has developed a package which has demonstrated the capability to efficiently monitor tens of thousands of entities on thousands of hosts, and has been in operation for over 4 years. The project has met the majority of its initial requirements, and also the majority of the requirements discovered along the way. This paper will describe what worked, and what did not, in the first 4 years of the NGOP Project at Fermilab; and we hope will provide valuable lessons for others considering undertaking even larger (GRID-scale) monitoring projects.

Distributed Computing Services / 253

Experiences with Data Indexing services supported by the NorduGrid middleware

A. Konstantinov¹ ; A. Waananen² ; B. Konya³ ; F. Ould-Saada⁴ ; J. Renner Hansen² ; J.L. Nielsen² ; M. Ellert⁵ ; O. Smirnova³ ; P. Eerola³ ; T. Ekelof⁶

¹ *Vilnius University, Lithuania*

² *Niels Bohr Institute, Copenhagen, Denmark*

³ *Lund University, Sweden*

⁴ *University of Oslo, Norway*

⁵ *CERN*

⁶ *Uppsala University, Sweden*

Auteur(s) contact: oxana.smirnova@hep.lu.se, langgard@nbi.dk

The NorduGrid middleware, ARC, has integrated support for querying and registering to Data Indexing services such as the Globus Replica Catalog and Globus Replica Location Server. This support allows one to use these Data Indexing services for for example brokering during job-submission, automatic registration of files and many other things. This integrated support is complemented by a set of command-line tools for registering to and querying these Data Indexing services.

In this talk we will describe experiences with these Data Indexing services both from a daily work point of view and in production environments such as the Atlas Data-Challenges 1 and 2. We will describe the advantages of such Data Indexing services as well as their shortcomings. Finally we will present a proposal for an extended Data Indexing service which should deal with the shortcomings described. The development of such a Data Indexing service is being planned at the moment.

Distributed Computing Services / 305

Experiences with the gLite Grid Middleware

A. Demichev¹; A. Maier²; A. Peters²; B. Koblitz²; D. Feichtinger³; D. Liko²; F. Orellana²; J. Herralá²; J. Moscicki²; M. Lamanna⁴; T. Chen⁵; V. Pose⁶; W. Ueng⁵

¹ *Moscow State University*

² *CERN*

³ *Swiss Institute of Particle Physics*

⁴ *CERN-IT*

⁵ *Academica Sinica, Nanking*

⁶ *JINR, Dubna*

The ARDA project was started in April 2004 to support the four LHC experiments (ALICE, ATLAS, CMS and LHCb) in the implementation of individual production and analysis environments based on the EGEE middleware.

The main goal of the project is to allow a fast feedback between the experiment and the middleware development teams via the construction and the usage of end-to-end prototypes allowing users to perform analyses out of the present data sets from recent montecarlo productions.

The LCG ARDA project is contributing to the development of the new EGEE Grid middleware by exercising it with realistic analysis systems developed within the four LHC experiments. We will present our experiences in using the EGEE middleware in first prototypes developed by the experiments together with the ARDA project. We will cover aspects such as the usability of individual components of the middleware and give an overview on which components are used by which experiments.

Poster Session 2 - Tableau: 53 / 116

Experiment Software Installation experience in LCG-2

A. Delgado Peris¹; A. Sciaba¹; F. Donno¹; P. Mendez Lorenzo²; R. santinelli¹; S. Campana¹

¹ *CERN/IT/GD*

² *CERN IT/GD*

Auteur(s) contact: roberto.santinelli@cern.ch

The management of Application and Experiment Software represents a very common issue in emerging grid-aware computing infrastructures.

While the middleware is often installed by system administrators at a site via customized tools that serve also for the centralized management of the entire computing facility, the problem of installing, configuring and validating Gigabytes of Virtual Organization (VO) specific software or frequently changing user applications remains an open issue.

Following the requirements imposed by the experiments, in the LHC Computing Grid (LCG) Experiment Software Managers (ESM) are designated people with privileges of installing, removing and validating software for a specific VO on a per site basis.

They can manage univocally identifying tags in the LCG Information System to announce the availability of a specific software version.

Users of a VO can then select, via the published tag, sites to run their jobs.

The solution adopted by LCG has mainly served its purpose but it presents many problems.

The requirement imposed by the present solution for the existence of a shared file-system in a computing farm poses performance, reliability and scalability issues for large installations.

With this work we present a more flexible service based on P2P technology that has been designed to tackle the limitation of the current system. This service allows the ESM to propagate the installation occurring in a given WN to the rest of the farm elements. We illustrate the deployment, the design, preliminary results obtained and the feedback from the LHC experiments and sites that have adopted it.

Poster Session 3 - Tableau: 65 / 322

Extending EGS with SVG for Track Visualization

B. White¹ ; K. Shrivastava² ; R. Cowan³

¹ STANFORD LINEAR ACCELERATOR CENTER (SLAC)

² Stanford Linear Accelerator Center/Stanford University

³ Stanford Linear Accelerator Center

Auteur(s) contact: bebo@slac.stanford.edu

The Electron Gamma Shower (EGS) Code System at SLAC is designed to simulate the flow of electrons, positrons and photons through matter at a wide range of energies. It has a large user base among the high-energy physics community and is often used as a teaching tool through a Web interface that allows program input and output. Our work aims to improve the user interaction and shower visualization model of the EGS Web interface. Currently, manipulation of the graphical output (a GIF file) is limited to simple operations like panning and zooming, and each such operation requires server-side calculations. We use SVG (Scalable Vector Graphics) to allow a much richer set of operations, letting users select a track and visualize it with the aid of 3-D rotations, adjustable particle display intensities, and interactive display of the interactions happening over time. A considerable advantage of our method is that once a track is selected for visualization, all further manipulations on that track can be done client-side without requiring server-side calculations. We hence combine the advantages of the SVG format (powerful interaction models over the Web) with those of conventional image formats (file size independent of scene complexity) to allow a composite set of operations for users, and enhance the value of EGS as a pedagogical tool.

Event Processing / 143

FAMOS, a FAsT MOnTe-Carlo Simulation for CMS

F. Beaudette¹

¹ CERN

Auteur(s) contact: florian.beaudette@cern.ch

An object-oriented FAsT MOnTe-Carlo Simulation (FAMOS) has recently been developed for CMS to allow rapid analyses of all final states envisioned at the LHC while keeping a high degree of accuracy for the detector material description and the related particle interactions. For example, the simulation of the material effects in the tracker layers includes charged particle energy loss by ionization and multiple scattering, electron Bremsstrahlung and photon conversion. The particle showers are developed in the calorimeters with an emulation of GFLASH, finely interfaced with the calorimeter geometry (e.g., crystal positions, cracks, rear leakage, etc).

As the same software framework is used for FAMOS and ORCA (the full Object-oriented Reconstruction software for CMS Analysis), the various Physics Objects (electrons, photons, muons, taus, jets, missing ET, charged particle tracks, ...) can be accessed with a similar code with both fast and full simulation, thus allowing any analysis algorithm to be transported from FAMOS to ORCA (and later, to data analysis and DST reading) or vice-versa without any additional work.

Altogether, a gain in CPU time of about a hundred can be achieved with respect to the full simulation, with little loss in precision.

Event Processing / 229

Fast reconstruction of tracks in the inner tracker of the CBM experiment

I. Kisel¹

¹ UNIVERSITY OF HEIDELBERG, KIRCHHOFF INSTITUTE OF PHYSICS

Typical central Au-Au collision in the CBM experiment (GSI, Germany) will produce up to 700 tracks in the inner tracker. Large track multiplicity together with presence of nonhomogeneous magnetic field make reconstruction of events complicated.

A cellular automaton method is used to reconstruct tracks in the inner tracker. The cellular automaton algorithm creates short track segments in neighbored detector planes and links them into tracks. Being essentially local and parallel the cellular automaton avoids exhaustive combinatorial search, even when implemented on conventional computers. Since the cellular automaton operates with highly structured information, the amount of data to be processed in the course of the track search is significantly reduced. The method employs a very simple track model which leads to utmost computational simplicity and fast algorithm.

Efficiency of track reconstruction for particles detected in at least three stations is presented. Tracks of high momentum particles are reconstructed very well with efficiency about 98%, while multiple scattering in detector material leads to lower reconstruction efficiency of slow particles.

Event Processing / 173

Fast tracking for the ATLAS LVL2 Trigger

D. Emelianov¹ ; F. PARODI² ; H. Drevermann³ ; J. Baines¹ ; M. Sutton⁴ ; N. Konstantinidis⁵

¹ RAL

² INFN Genova, Italy

³ CERN

⁴ UCL

⁵ UNIVERSITY COLLEGE LONDON

Auteur(s) contact: nikolaos.konstantinidis@cern.ch

We present a set of algorithms for fast pattern recognition and track reconstruction using 3D space points aimed for the High Level Triggers (HLT) of multi-collision hadron collider environments. At the LHC there are several interactions per bunch crossing separated

along the beam direction, z . The strategy we follow is to (a) identify the z -position of the interesting interaction prior to any track reconstruction; (b) select groups of space points pointing back to this z -position, using a histogramming technique which avoids performing any combinatorics; and (c) proceed to the combinatorial tracking only within the individual groups of space points. The validity of this strategy will be demonstrated with results in terms of timing and physics performance for the LVL2 trigger of ATLAS at the LHC, although the strategy is generic and can be applied to any multi-collision hadron collider experiment.

In addition, the algorithms are conceptually simple, flexible and robust and hence appropriate for use in demanding, online environments. We will also make qualitative comparisons with an alternative, complimentary strategy, based on the use of look-up tables for handling combinatorics, that has been developed for the ATLAS LVL2 trigger. These algorithms have been used for the results that appear in the ATLAS HLT, DAQ and Controls Technical Design Report, which was recently approved by the LHC Committee.

Online Computing / 432

Fault Tolerance and Fault Adaption for High Performance Large Scale Embedded Systems

P. Sheldon¹

¹ VANDERBILT UNIVERSITY

Auteur(s) contact: paul.sheldon@vanderbilt.edu

The BTeV experiment, a proton/antiproton collider experiment at the Fermi National Accelerator Laboratory, will have a trigger that will perform complex computations (to reconstruct vertices, for example) on every collision (as opposed to the more traditional approach of employing a first level hardware based trigger). This trigger requires large-scale fault adaptive embedded software: with thousands of processors involved in performing event filtering in the trigger farm fault conditions must be given proper treatment. Without fault mitigation, it is conceivable that the trigger system will experience failures at a high enough rate to have an unacceptable negative impact on BTeV's physics goals. The RTES (Real Time Embedded Systems) collaboration is a group of physicists, engineers, and computer scientists working to address the problem of reliability in large-scale clusters with real-time constraints such as this. Resulting infrastructure must be highly scalable, verifiable, extensible by users, and dynamically changeable. An initial prototype has been built to test design ideas and methods for the final system, and a larger scale and more ambitious prototype is currently under construction. I will discuss the lessons learned from these prototypes as well as the overall design and deliverables for the BTeV experiment.

Poster Session 2 - Tableau: 67 / 55

Federating Grids: LCG meets Canadian HEPGrid

A. Agarwal¹ ; A. Dimopoulos¹ ; B. Caron² ; C. Lindsay¹ ; D. Vanderster¹ ; G. Mateescu³ ; L. Klektau¹ ; M. Vetterli⁴ ; R. Impey³ ; R. Sobie¹ ; R. Walker⁵

¹ University of Victoria

² *University of Alberta/TRIUMF*

³ *Ottawa*

⁴ *Simon Fraser University/TRIUMF*

⁵ *Simon Fraser University*

A large number of Grids have been developed, motivated by geo-political or application requirements. Despite being mostly based on the same underlying middleware, the Globus Toolkit, they are generally not inter-operable for a variety of reasons. We present a method of federating those disparate grids which are based on the Globus Toolkit, together with a concrete example of interfacing the LHC grid(LCG) with HEPGrid. HEPGrid consists of shared resources, at several Canadian research institutes, which are exposed via Globus gatekeepers, and makes use of Condor-G for resource advertisement, matchmaking and job submission. An LCG Computing Element(CE) based at the TRIUMF Laboratory hosts a HEPGrid User Interface(UI) which is contained within a custom jobmanager. This jobmanager appears in the LCG information system as a normal CE publishing an aggregation of the HEPGrid resources. The interface interprets the incoming job in terms of HEPGrid UI usage, submits it onto HEPGrid, and implements the jobmanager 'poll' and 'remove' methods, thus enabling monitoring and control across the grids. In this way non-LCG resources are integrated into LCG, without the need for LCG middleware on those resources. The same method can be used to create interfaces between other grids, with the details of the child-Grid being fully abstracted into the interface layer. The LCG-HEPGrid interface is operational, and has been used to federate 1300 CPU's at 4 sites into LCG for the Atlas Data Challenge (DC2).

Distributed Computing Services / 392

File-Metadata Management System for the LHCb Experiment

C. CIOFFI¹ ; F. LOVERRE² ; J. CLOSIER² ; M. FRANK² ; S. PONCE²

¹ *Oxford University*

² *CERN*

Auteur(s) contact: c.cioffi1@physics.ox.ac.uk

The LHCb experiment needs to store all the information about the datasets and their processing history of recorded data resulting from particle collisions at the LHC collider at CERN as well as of simulated data.

To achieve this functionality a design based on data warehousing techniques was chosen, where several user-services can be implemented and optimized individually without losing functionality nor performance. This approach results in an experiment-independent and flexible system. It allows fast access to the catalogue of available data, to detailed history information and to the catalogue of data replicas. Queries can be made based on these three sets of information. A flexible underlying database schema allows the implementation and evolution of these services without the need to change the basic database schema. The consequent implementation of interfaces based on XML-RPC allows to access and to modify the stored information using a well defined encapsulating API.

Distributed Computing Services / 204

FroNtier: High Performance Database Access Using Standard Web Components in a Scalable Multi-tier Architecture

B. Blumenfeld¹ ; D. Litvintsev² ; J. Kowalkowski³ ; L. Lueking⁴ ; M. Mathis¹ ; M. Paterno³ ; P. Maksimovic¹ ; S. Kosyakov³ ; S.P. White³

¹ *John Hopkins University*

² *FNAL*

³ *Fermilab*

⁴ *FERMILAB*

Auteur(s) contact: lueking@fnal.gov

A high performance system has been assembled using standard web components to deliver database information to a large number (thousands?) of broadly distributed clients. The CDF Experiment at Fermilab is building processing centers around the world imposing a high demand load on their database repository. For delivering read-only data, such as calibrations, trigger information and run conditions data, we have abstracted the interface that clients use to retrieve database objects. A middle tier is deployed that translates client requests into database specific queries and returns the data to the client as HTTP datagrams. The database connection management, request translation, and data encoding are accomplished in servlets running under Tomcat. Squid Proxy caching layers are deployed near the Tomcat servers as well as close to the clients to significantly reduce the load on the database and provide a scalable deployment model. This system is highly scalable, readily deployable, and has a very low administrative overhead for data delivery to a large, distributed audience. Details of how the system is built and used will be presented including its architecture, design, interfaces, administration, and performance measurements.

Event Processing / 323

From Geant 3 to Virtual Monte Carlo: Approach and Experience

J. Lauret¹ ; M. POTEKHIN² ; V. Perevoztchikov² ; Y. FISYAK²

¹ *Brookhaven National Laboratory*

² *BROOKHAVEN NATIONAL LABORATORY*

The STAR Collaboration is currently using simulation software based on Geant 3. The emergence of the new Monte Carlo simulation packages, coupled with evolution of both STAR detector and its software, requires a drastic change of the simulation framework.

We see the Virtual Monte Carlo (VMC) approach as providing a layer of abstraction that facilitates such transition.

The VMC platform is a candidate to replace the present legacy software, and help avoid its certain shortcomings, such as the use of a particular algorithmic language to describe the detector geometry. It will also allow us to introduce a more flexible in-memory representation of the geometry.

The Virtual Monte Carlo concept includes a platform-neutral kernel of the application, to the highest degree possible.

This kernel is then equipped with interfaces to the modules responsible for simulating the physics of particle propagation, and tracking.

We consider the geometry description classes in the ROOT system (in its latest form known as TGeo classes) as a good choice for the in-memory geometry representation.

We present an application design based on the Virtual Monte Carlo, along with the results of testing, benchmarking and comparison to Geant 3. Internal event representation and IO model will be also discussed.

Event Processing / 405

Full Event Reconstruction in Java

N. Graf¹

¹ SLAC

Auteur(s) contact: norman.graf@slac.stanford.edu

We describe a Java toolkit for full event reconstruction and analysis. The toolkit is currently being used for detector design and physics analysis for a future linear e+ e- linear collider. The components are fully modular and are available for tasks from digitization of tracking detector signals through to cluster finding, pattern recognition, fitting, jetfinding, and analysis. We discuss the architecture as well as the implementation for several candidate detector designs.

Poster Session 1 - Tableau: 13 / 237

Future processors: What is on the horizon for HEP farms?

S. Jarp¹

¹ CERN

Auteur(s) contact: sverre.jarp@cern.ch

In 1995 I predicted that the dual-processor PC would start invading HEP computing and a couple of years later the x86-based PC was omnipresent in our computing facilities. Today, we cannot imagine HEP computing without thousands of PCs at the heart. This talk will look at some of the reasons why we may one day be forced to leave this sweet-spot. This would be not because we (the HEP community) want to, but rather because other market forces may pull in different directions. Amongst such forces, I will review the new generation of powerful game consoles where IBM's Power processor is currently making strong inroads. Then I will look at the huge mobile market where low-powered processing rules rather than power-hungry DP Xeon/Xeon-like processors, and thirdly I will explore in my talk the promise of enterprise servers with a large number of processors on each die (so-called Core Multi-Processors). For all the scenarios, we must, of course, keep in mind that HEP can only move when the price-performance ratio is right.

Grid Security / 194

G-PBox: a Policy Framework for Grid Environments

A. Caltroni¹ ; A. Ferraro² ; A. Ghiselli³ ; G. Rubini² ; R. Zappi² ; V. Ciaschini²

¹ *INFN PD*

² *INFN CNAF*

³ *INFN-CNAF*

Auteur(s) contact: vincenzo.ciaschini@cnaif.infn.it

A key feature of Grid systems is the sharing of its resources among multiple Virtual Organizations (VOs). The sharing process needs a policy framework to manage the resource access and usage. Generally Policy frameworks exist for farms or local systems only, but now, for Grid environments, a general, and distributed policy system is necessary.

Generally VOs and local systems have contracts that regulate the resource usage, hence complex relationships among these entities implying different kind of policies may exist: VOs oriented, local systems oriented, and a mix of these ones. We propose an approach to the representation, and management of such policies: the Grid Policy Box (G-PBox) framework. The approach is based on a set of databases belonging hierarchically-organised levels distributed onto the Grid and VOs structures. Each level contains only policies regarding itself. These levels have to communicate among themselves to accommodate for mixed policies, originating the need for a secure communication service framework, - for privacy reasons, - with the ability to sort and dispatch various kind of policies to the involved parties.

In this paper we present our first implementation of the G-PBox, and its architecture details, and we discuss the plans for G-PBox-related application and research.

Poster Session 2 - Tableau: 9 / 380

GILDA: a Grid for dissemination activities

A. Carrieri¹ ; A. Falzone² ; A. Italiano³ ; E. Ferro⁴ ; E. Giorgio¹ ; G. Andronico¹ ; G. La Rocca¹ ; G. Platania¹ ; M. Pappalardo¹ ; M. Verlatto⁴ ; R. Barbera⁵ ; R. Catania¹ ; S. Monforte¹ ; V. Ardizzone¹

¹ *INFN Catania*

² *NICE s.r.l.*

³ *INFN CNAF*

⁴ *INFN Padova*

⁵ *Univ. Catania and INFN Catania*

Auteur(s) contact: roberto.barbera@ct.infn.it

Computational and data grids are now entering a more mature phase where experimental test-beds are turned into production quality infrastructures operating around the clock. All this is becoming true both at national level, where an example is the Italian INFN production grid (<http://grid-it.cnaif.infn.it>), and at the continental level, where the most striking example is the European Union EGEE Project Infrastructure (<http://www.eu-egee.org>).

However, the impact of grid technologies on the next future way of doing e-science and research in Europe will be proportional to the capability of National and European Grid Infrastructures to attract and serve many diverse scientific and industrial communities through serious and detailed dissemination and tutoring programs.

In this contribution we present GILDA, the Grid Infn Laboratory for Dissemination Activities (<http://gilda.ct.infn.it>). GILDA is a complete suite of grid elements

(Certification Authority, Virtual Organization, Distributed Test-bed, Grid Demonstrator, etc.) completely devoted to dissemination activities. GILDA can also act as a fast-prototyping test-bed where to start the porting/interfaces of new applications with the grid middle-ware. The use and exploitation of GILDA in the context of the Network Activities of the EGEE Project will be discussed.

Poster Session 2 - Tableau: 59 / 148

GROSS: an end user tool for carrying out batch analysis of CMS data on the LCG-2 Grid.

B. MacEvoy¹ ; D. Colling¹ ; H. Tallini¹ ; W. Stuart¹

¹ IMPERIAL COLLEGE LONDON

Auteur(s) contact: h.tallini@imperial.ac.uk, stuart.wakefield@imperial.ac.uk

GROSS (GRidified Orca Submission System) has been developed to provide CMS end users with a single interface for running batch analysis tasks over the LCG-2 Grid. The main purpose of the tool is to carry out job splitting, preparation, submission, monitoring and archiving in a transparent way which is simple to use for the end user. Central to its design has been the requirement for allowing multi-user analyses, and to accomplish this all persistent information is stored on a backend MySQL database. This database is additionally shared with BOSS, to which GROSS interfaces in order to provide job submission and real time monitoring capability.

In this paper we present an overview of GROSS's architecture and functionality and report on first user tests of the system using CMS Data Challenge 2004 data (DC04).

Poster Session 3 - Tableau: 20 / 376

Geant4 as Simulation Toolkit addressed to interplanetary manned missions studies: required developments and improvements

M.G. Pia¹ ; S. Guatelli¹

¹ INFN Genova, Italy

Auteur(s) contact: guatelli@ge.infn.it

The study of the effects of space radiation on astronauts in an important concern of space missions for the exploration of the Solar System. The radiation hazard to crew is critical to the feasibility of interplanetary manned missions.

To protect the crew, shielding must be designed, the environment must be anticipated and monitored, and a warning system must be put in place.

A Geant4 simulation has been developed for a preliminary quantitative study of vehicle concepts and Moon surface habitat designs, and the radiation exposure of crews therein. This project is defined in the context of the European AURORA programme, whose primary object is to study solutions for the robotic and human exploration of the Solar System, with Mars, the Moon and the asteroids as the most likely objects.

This study intends to evaluate whether the energy range typical of the radiation environment of interplanetary missions is adequately treated in Geant4 physics

packages, for all the major types of particles involved, identifying the availability of appropriate electromagnetic and hadronic physics models and verifying the status of their validation. Recommendations for further Geant4 developments or improvements and of further validation tests, necessary for the interplanetary manned missions, are issued as a result of this study.

Event Processing / 314

Geant4: status and recent developments

H.P. wellisch¹ ; J. Apostolakis¹ ; M. Asai²

¹ CERN

² SLAC

Auteur(s) contact: john.apostolakis@cern.ch

Geant4 is relied upon in production for increasing number of HEP experiments and for applications in several other fields. Its capabilities continue to be extended, as its performance and modelling are enhanced.

This presentation will give an overview of recent developments in diverse areas of the toolkit. These will include, amongst others, the optimisation for complex setups using different production thresholds, improvements in the propagation in fields, and highlights from the physics processes and event biasing.

In addition it will note the physics validation effort undertaken in collaboration with a number of experiments, groups and users.

Poster Session 2 - Tableau: 23 / 480

Generic logging layer for the distributed computing

G. Van Buren¹ ; J. Lauret² ; V. Fine¹

¹ BROOKHAVEN NATIONAL LABORATORY

² Brookhaven National Laboratory

Most HENP experiment software includes a logging or tracing API allowing for displaying in a particular format important feedback coming from the core application. However, inserting log statements into the code is a low-tech method for tracing the program execution flow and often leads to a flood of messages in which the relevant ones are occluded. In a distributed computing environment, accessing the information via a log-file is no longer applicable and the approach fails to provide runtime tracing.

Running a job involves a chain of events where many components are involved often written in diverse languages and not offering a consistent and easily adaptable interface for logging important events.

We will present an approach based on a new generic layer built on top of a logger family derived from the Jakarta log4j project that includes log4cxx, log4c, log4perl packages. This provides consistency across packages and framework.

Additionally, the power of using log4j, is the possibility to enable logging (or features) at runtime without modifying the application binary or the wrapper layers.

We provide a C++ abstract class library that serves as a proxy between the

application framework and the distributed environment. The approach is designed so that the debugging statements can remain in shipped code without incurring a heavy performance cost. Logging equips the developer with as detailed context as necessary for application failures, from testing, quality assurance to a production mode limited amount of information. We will explain and show its implementation in the STAR production environment.

Event Processing / 49

Genetic Programming and its application to HEP

E. Vaandering¹

¹ VANDERBILT UNIVERSITY

Auteur(s) contact: ewv@fnal.gov

Genetic programming is a machine learning technique, popularized by Koza in 1992, in which computer programs which solve user-posed problems are automatically discovered. Populations of programs are evaluated for their fitness of solving a particular problem. New populations of ever increasing fitness are generated by mimicking the biological processes underlying evolution. These processes are principally genetic recombination, mutation, and survival of the fittest.

Genetic programming has potential advantages over other machine learning techniques such as neural networks and genetic algorithms in that the form of the solution is not specified in advance and the program can grow as large as necessary to adequately solve the posed problem.

This talk will give an overview and demonstration of the genetic programming technique and show a successful application in high energy physics: the automatic construction of an event filter for FOCUS which is more powerful than the experiment's usual methods of event selection. We have applied this method to the study of doubly Cabibbo suppressed decays of charmed hadrons (D^+ , D_s^+ , and Λ_c^+).

Computer Fabrics / 458

Gfarm v2: A Grid file system that supports high-performance distributed and parallel data computing

N. Soda¹ ; O. Tatebe² ; S. Matsuoka³ ; S. Sekiguchi⁴ ; Y. Morita⁵

¹ SRA

² GRID TECHNOLOGY RESEARCH CENTER, AIST

³ Tokyo Institute of Technology/National Institute of Informatics

⁴ Grid Technology Research Center, AIST

⁵ KEK

Auteur(s) contact: o.tatebe@aist.go.jp

Gfarm v2 is designed for facilitating reliable file sharing and high-performance distributed and parallel data computing in a Grid across administrative domains by providing a Grid file system. A

Grid

file system is a virtual file system that federates multiple file systems. It is possible to share files or data by mounting the virtual file system. This paper discusses the design and implementation of secure, robust, scalable and high-performance Grid file system.

The most time-consuming, but also the most typical, task in data computing such as high energy physics, astronomy, space exploration, human genome analysis, is to process a set of files in the same way. Such a process can be typically performed independently on every file in parallel, or at least have good locality. Gfarm v2 supports high-performance distributed and parallel computing for such a process by introducing a “Gfarm file”, a new “file-affinity” process scheduling based on file locations, and new parallel file access semantics. An arbitrary group of files possibly dispersed across administrative domains can be managed as a single Gfarm file. Each member file will be accessed in parallel in a new file view called “local file view” by a parallel process possibly allocated by file-affinity scheduling based on replica locations of the member files. File-affinity scheduling and new file view enable the owner computes” strategy, or move the computation to data” approach for parallel and distributed data computing of member files of a Gfarm file in a single system image.

Distributed Computing Services / 430

Global Distributed Parallel Analysis using PROOF and AliEn

A. Peters¹ ; F. Carminati¹ ; F. Rademakers¹ ; M. Ballintijn² ; P. BUNCIC¹ ; P. Canal³ ; R. Brun¹

¹ CERN

² MIT

³ FNAL

Auteur(s) contact: federico.carminati@cern.ch

The ALICE experiment and the ROOT team have developed a Grid-enabled version of PROOF that allows efficient parallel processing of large and distributed data samples. This system has been integrated with the ALICE-developed AliEn middleware. Parallelism is implemented at the level of each local cluster for efficient processing and at the Grid level, for optimal workload management of distributed resources. This system allows harnessing large Computing on Demand capacity during an interactive session. Remote parallel computations are spawned close to the data, minimising network traffic. If several copies of the data are available, a workload management system decides automatically where to send the task. Results are automatically merged and displayed at the user workstation. The talk will describe the different components of the system (PROOF, the parallel ROOT engine, and the AliEn middleware), the present status and future plans for the development and deployment and the consequences for the ALICE computing model.

Poster Session 2 - Tableau: 6 / 504

Global Grid User Support for LCG

G. GREIN^{None} ; H. DRES^{None} ; T. ANTONI¹ ; W. THOENE^{None}

¹ GGUS

For very large projects like the LHC Computing Grid Project (LCG) involving 8,000 scientists from all around the world, it is an indispensable requirement to have a well organized user support. The Institute for Scientific Computing at the Forschungszentrum Karlsruhe started implementing a Global Grid User Support (GGUS) after official assignment of the Grid Deployment Board in March 2003. For this purpose a web portal and a helpdesk application have been developed. As a single entry point for all Grid related issues and problems GGUS follows the objectives of providing news, documentation and status information about Grid resources. The user will find forms to submit and track service requests. GGUS collaborates with different support teams in the Grid environment like the Grid Operations Center and the Experiment Specific Support. They can access the helpdesk system via web interface. GGUS stores all the incoming trouble tickets and outgoing solutions in a central database and plans to build up a knowledge base where all the information can be offered in a structured manner.

As a prototype GGUS started operation at the Forschungszentrum Karlsruhe in October 2003 and supported local user groups of the German Tier 1 Computing Center, called GridKa. 4 month later the GGUS system was opened for the LCG community. The GGUS system will be explained and demonstrated. The present status of GGUS within the LCG environment will be discussed.

Distributed Computing Systems and Experiences / 366

Globally Distributed User Analysis Computing at CDF

A. Fella¹ ; A. Kreymer² ; A. Sill³ ; D. Han⁴ ; E. Lipeles⁵ ; F. Ratnikov⁶ ; F. Wuerthwein⁷ ; H-T. Lung⁸ ; H. Matsunaga⁹ ; H. Park¹⁰ ; I. Sfiligoi¹¹ ; K. Cho⁴ ; L. Groer¹² ; M. Burgon-Lyon¹³ ; M. Neubauer⁷ ; R. St.Denis¹³ ; R. Tafirout¹² ; S. Belforte¹⁴ ; S. Hou⁸ ; S. Stonjek¹⁵ ; V. Bartsch¹⁵

¹ INFN, Pisa

² Fermi National Accelerator Laboratory

³ TEXAS TECH UNIVERSITY

⁴ Kyungpook National University

⁵ UNIVERSITY OF CALIFORNIA SAN DIEGO

⁶ Rutgers University

⁷ University of California, San Diego

⁸ Academia Sinica, Taiwan

⁹ University of Tsukuba

¹⁰ Kyungpook National University/KISTI

¹¹ INFN, Frascati

¹² University of Toronto

¹³ University of Glasgow

¹⁴ INFN, Trieste

¹⁵ University of Oxford

To maximize the physics potential of the data currently being taken, the CDF collaboration at Fermi National

Accelerator Laboratory has started to deploy user analysis computing facilities at several locations throughout the world. Over 600 users are signed up and able to submit their physics analysis and simulation applications directly from their desktop or laptop computers to these facilities. These resources consist of a mix of customized computing centers and a decentralized version of our Central Analysis Facility (CAF) initially used at Fermilab, which we have designated Decentralized CDF Analysis Facilities (DCAFs).

We report on experience gained during the initial deployment and use of these resources for the summer conference season 2004. During this period, we allowed MC generation as well as data analysis of selected data samples at several globally distributed centers. In addition, we discuss a migration path from this first generation distributed computing infrastructure towards a more open implementation that will be interoperable with LCG, OSG and other general-purpose grid installations at the participating sites.

Core Software / 418

Go4 analysis design

Auteur(s): H. Essel¹

Co-auteur(s) D. Bertini¹ ; J. Adamczewski¹ ; M. Al-Turany¹ ; S. Linev¹

¹ GSI

Auteur(s) contact: h.essel@gsi.de

The GSI online-offline analysis system Go4 is a ROOT based framework for medium energy ion- and nuclear physics experiments. Its main features are a multithreaded online mode with a non-blocking Qt GUI, and abstract user interface classes to set up the analysis process itself which is organised as a list of subsequent analysis steps. Each step has its own event objects and a processor instance. It can handle its event i/o independently. It can be set up by macros or by generic a GUI. With respect to the more complex experiments planned at GSI, a configurable network of steps is required. Multiple IO channels per step and multiple references to steps can be set up by macros or via generic GUI. The required mechanisms are provided by an upgrade of the Go4 analysis step manager using the new ROOT TTasks. Support for IO configuration and references across the task tree is provided.

Poster Session 2 - Tableau: 50 / 267

GoToGrid - A Web-Oriented Tool in Support to Sites for LCG Installations

A. Retico¹ ; A. Usai¹ ; G. Diez-Andino Sancho¹ ; I. tkachev¹ ; M. Schulz¹ ; O. Keeble¹ ; P. Nyczyk¹

¹ CERN

Auteur(s) contact: antonio.retico@cern.ch

The installation and configuration of LCG middleware, as it is currently being done, is complex and delicate.

An “accurate” configuration of all the services of LCG middleware requires a deep knowledge of the inside dynamics and hundreds of parameters to be dealt with. On the other hand, the number of parameters and flags that are strictly needed in order to run a working “default” configuration of the middleware is relatively small, due to the fact that the values to be set mainly deal with environment configuration and with a limited set of possible operation scenarios.

This “default” configuration appears to be the most suitable for sites joining LCG for the first time.

The GoToGrid system is aimed to support Site Administrators to easily perform such a configuration.

G2G combines the gathering of configuration information, provided by sites, with the dynamic adaptive creation of customized documentation and installation tools.

By using a web interface and being requested only for the relevant configuration information, site Administrators will be able to design the desired configuration of their own LCG site.

Site configuration data is collected and stored in a well defined format liable to be used as the interface to different configuration management tools.

Poster Session 3 - Tableau: 21 / 425

GraXML

J. Hrivnac¹

¹ LAL

Auteur(s) contact: julius.hrivnac@cern.ch

GraXML is the framework for manipulation and visualization of 3D geometrical objects in space. The full framework consists of the GraXML toolkit, libraries implementing Generic and Geometric Models and end-user interactive front-ends. GraXML Toolkit provides a foundation for operations on 3D objects (both detector elements and events). Each external source of 3D data is automatically translated into Generic Model which is then analyzed and translated into Geometric Model using GraXML modules. The construction of this Geometric Model is parametrised by several parameters (optimization level, quality level, ...) so that it can be used in applications with different requirements (graphical or not). Two visualization applications are provided in the GraXML framework: GraXML Interactive Display and GraXML Converter into various 3D geometry formats. Other applications can be easily developed.

The presentation will concentrate on GraXML graphical capabilities and relation with geometric data providers. The difference between specific GraXML features and properties of other similar tools will be highlighted. The questions of different visualization needs and possibilities for different kinds of geometrical data will be also explained.

Plenary / 526

Grand Challenges facing Storage Systems

J. MENON^{None}

In this talk, we will discuss the future of storage systems. In particular, we will focus on several big challenges which we are facing in storage, such as being able to build, manage and backup really massive storage systems, being able to find information of interest, being able to do long-term archival of data, and so on. We also present ideas and research being done to address these challenges, and provide a perspective on how we expect these challenges to be resolved as we go forward.

Distributed Computing Systems and Experiences / 319**Grid Collector: Using an Event Catalog to Speed up User Analysis in Distributed Environment**

A. Shoshani¹ ; J. Lauret² ; K. Wu³ ; V. Perevoztchikov⁴ ; W-M. Zhang⁵

¹ *Lawrence Berkeley National Lab*

² *Brookhaven National Laboratory*

³ *LAWRENCE BERKELEY NATIONAL LAB*

⁴ *Brookhaven National Lab*

⁵ *Kent State University*

Auteur(s) contact: kwu@lbl.gov

Nuclear and High Energy Physics experiments such as STAR at BNL are generating millions of files with PetaBytes of data each year. In most cases, analysis programs have to read all events in a file in order to find the interesting ones.

Since most analyses are only interested in some subsets of events in a number of files, a significant portion of the computer time is wasted on reading the unwanted events. To address this issue, we developed a software system called the Grid Collector. The core of the Grid Collector is an "Event Catalog".

This catalog can be efficiently searched with compressed bitmap indices. Tests show that it can index and search STAR event data much faster than database systems.

It is fully integrated with an existing analysis framework so that a minimal effort is required to use the Grid Collector in an analysis program. In addition, by taking advantage of existing file catalogs, Storage Resource Managers (SRMs) and GridFTP, the Grid Collector automatically downloads the needed files anywhere on the Grid without user intervention.

The Grid Collector can significantly improve user productivity. The improvement in productivity is more significant as users converge toward searching for rare events, because only the rare events are read into memory and the necessary files are automatically located and downloaded through the best available route. For a user that typically performs computation on 50% of the events, using the Grid Collector could reduce the turn around time by a half.

Poster Session 2 - Tableau: 22 / 126

Grid Deployment Experiences: The path to a production quality LDAP based grid information system

L. Field¹ ; M. Schulz¹

¹ *CERN*

Auteur(s) contact: laurence.field@cern.ch

This paper reports on the deployment experience of the defacto grid information system, Globus MDS, in a large scale production grid. The results of this experience led to the development of an information caching system based on a standard openLDAP database. The paper then

describes how this caching system was developed further into a production quality information system including a generic framework for information providers. This includes the deployment and operation experience and the results from performance tests on the information system to assess the scalability limits of it.

Distributed Computing Systems and Experiences / 182

Grid Enabled Analysis : Architecture, prototype and status

A. Anjum¹ ; C. Steenberg² ; D. Bourilkov³ ; F. van Lingen⁴ ; H. Newman⁵ ; I. LeGrand² ; J. Bunn² ; J. Uk In³ ; L. Chitnis³ ; M. Kulkarni⁶ ; M. Thomas² ; P. Avery³ ; R. Cavanaugh³ ; T. Azim¹

¹ *NUST, Pakistan*

² *California Institute of Technology*

³ *University of Florida*

⁴ *CALIFORNIA INSTITUTE OF TECHNOLOGY*

⁵ *Caltech*

⁶ *UNIVERSITY OF FLORIDA*

Auteur(s) contact: fvlingen@caltech.edu

In this paper we report on the implementation of an early prototype of distributed high-level services supporting grid-enabled data analysis within the LHC physics community as part of the ARDA project within the context of the GAE (Grid Analysis Environment) and begin to investigate the associated complex behaviour of such an end-to-end system. In particular, the prototype integrates a typical physics user interface client (ROOT), a uniform web-services interface to grid services (Clarens), a virtual data service (Chimera), a request scheduling service (Sphinx), a monitoring service (MonALISA), a workflow execution service (Virtual Data Toolkit Client), a remote data file service (Clarens), a grid resource service (Virtual Data Toolkit Server), a replica location service/meta data catalog (RLS/POOL), an analysis session management system (CAVES) and a fine grain monitor system for job submission (BOSS).

For testing and evaluation purposes, the prototype is deployed across a modest sized U.S. regional CMS Grid Test-bed (consisting of sites in California, Florida, Fermilab) and is in the early stages of exhibiting interactive remote data access demonstrating interactive workflow generation and collaborative data analysis using virtual data and data provenance, as well as showing non-trivial examples of policy based scheduling of requests in a resource constrained grid environment. In addition, the prototype is used to characterize the system performance as a whole, including the determination of request-response latencies in a distributed service model and the classification of high-level failure modes in a complex system.

Plenary / 531

Grid Security

D. KELSEY^{None}

The aim of Grid computing is to enable the easy and open sharing of resources between large and highly distributed communities of scientists and institutes across many independent administrative domains. Convincing site security officers and computer centre managers to allow this to happen in view of today's ever-increasing Internet security problems is a major challenge. Convincing users and application developers to take security seriously is equally difficult. This paper will describe the main Grid security issues, both in terms of technology and policy, that have been tackled over recent years in LCG and related Grid projects. Achievements to date will be described and opportunities for future improvements will be addressed.

Distributed Computing Services / 490

Grid2003 Monitoring, Metrics, and Grid Cataloging System

A. Zahn¹ ; B K. Kim² ; C. Prescott² ; I. Fisk³ ; I. Legrand⁴ ; J. Hicks⁵ ; J. Rodriguez² ; J. Weigand³ ; L. Grundhoefer⁵ ; M. Green⁶ ; M. Mambelli¹ ; P. Avery² ; R. Gardner¹ ; R. Quick⁵

¹ *University of Chicago*

² *UNIVERSITY OF FLORIDA*

³ *FERMI NATIONAL LABORATORY*

⁴ *CERN*

⁵ *Indiana University*

⁶ *University of Buffalo*

Auteur(s) contact: marco@hep.uchicago.edu, bockjoo@phys.ufl.edu

Grid computing involves the close coordination of many different sites which offer distinct computational and storage resources to the Grid user community. The resources at each site need to be monitored continuously. Static and dynamic site information need to be presented to the user community in a simple and efficient manner.

This paper will present both the design and implementation of the Grid3 monitoring infrastructure and the design details and the functionalities of a new application called the Gridcat.

The Grid3 monitoring architecture follows a user-oriented design that specifies standard metrics and uses different underlying monitoring tools to collect them and build a very diversified framework. In the monitoring framework we integrated existing tools, extended their functionality and developed original new tools. The main tools used include ACDC Job Monitoring from University of Buffalo, Ganglia, a preliminary version of Gridcat, Globus MDS, the University of Chicago Grid telemetry MDViewer, and US CMS MonALISA. From the collected data is extracted information of interest for the VOs participating in the Grid, for example resources provided and used by all VOs and jobs submitted by each VO.

The Gridcat shows site status using a web interface that is simple and powerful enough to determine the site's readiness to accept grid applications by collecting and storing dynamic site information to a database. The status information displayed by the prototype Gridcat was used extensively by the Grid2003 project as a coordination point for the grid operations center.

Plenary / 171

Grid3: An Application Grid Laboratory for Science

R. Gardner¹

¹ *University of Chicago*

Auteur(s) contact: robert.w.gardner@cern.ch

The U.S. Trillium Grid projects in collaboration with High Energy Experiment groups from the Large Hadron Collider (LHC), ATLAS and CMS, Fermi-Lab's BTeV, members of the LIGO, SDSS collaborations and groups from other scientific disciplines and computational centers have deployed a multi-VO, application-driven grid laboratory ("Grid3"). The grid laboratory has sustained for several months the production-level services required by the participating experiments. The deployed infrastructure has been operating since November 2003 with 27 sites, a peak of 2800 processors, work loads from 10 different applications exceeding 1300 simultaneous jobs, and data transfers among sites of greater than 2 TB/day.

The Grid3 infrastructure was deployed from grid level services provided by groups and applications within the collaboration. The services were organized into four distinct "grid level services" including: Grid3 Packaging, Monitoring and Information systems, User Authentication and the iGOC Grid Operations Center. In this paper we describe the Grid3 operational model, deployment strategies, and site installation and configuration procedures. We describe the grid middleware components used, how the components were packaged and deployed on sites each under its own local administrative domain, and how the pieces fit together to form the Grid3 grid infrastructure.

Plenary / 530

Griding The Nordic Supercomputing Infrastructure

B A. YNNERMAN^{None}

This talk gives a brief overview of recent development of high performance computing and Grid initiatives in the Nordic region. Emphasis will be placed on the technology and policy demands posed by the integration of general purpose supercomputing centers into Grid environments. Some of the early experiences of bridging national eBorders in the Nordic region will also be presented.

Rather than giving an exhaustive presentation of all projects in the Nordic countries the presentation uses selected examples of Grid projects to show the potential as well as some of the current limitations of Grids.

Plans for a common Nordic Grid Core Facility are currently being made. The presentation gives an overview of these plans and the status of the project. It will also cover a few examples of Nordic Grid initiatives in more detail such as the recently launched SweGrid test bed for production.

Core Software / 169

Guidelines for Developing a Good GUI

F. Rademakers¹; I. Antcheva¹; R. Brun¹

¹ *CERN*

Auteur(s) contact: ilka.antcheva@cern.ch

Designing a usable, visually-attractive GUI is somewhat more difficult than it appears at a first glance. The users, the GUI designers and the programmers are three important parts involved in this process and everyone has a comprehensive view on the aspects of the application goals, as well as the steps that have to be taken to meet successfully the application requirements. The fundamental GUI design principles and the main programming aspects are discussed in this paper.

Key topics include:

- User requirements: identifying users and support different user profiles - from beginners to advanced users
- Close relationship between the GUI widgets, user actions, tasks and user goals
- Task-analysis methods
- Prototypes development and testing prototypes
- General design considerations
- Effective GUI design keys, guidelines and style guides

Event Processing / 238

H100 - an analysis framework for H1

J. Katzy¹

¹ DESY, HAMBURG

Auteur(s) contact: judith.katzy@desy.de

During the years 2000 and 2001 the HERA machine and the H1 experiment performed substantial luminosity upgrades. To cope with the increased demands on data handling an effort was made to redesign and modernize the analysis software. Main goals were to lower turn-around time for physics analysis by providing a single framework for data storage, event selection, physics analysis and event display. The new object oriented analysis environment is using C++ and is based on the RooT framework. Data layers with a high level of abstraction are defined, i.e. physics particles, event summary information and user specific information.

A generic interface makes the use of reconstruction output stored in BOS format transparent to the user. Links between all data layers and partial event reading allow correlating quantities of different abstraction levels with high performance. Detailed physics analysis is performed by passing transient data between different analysis modules. Binding of existing fortran based libraries on demand allows the use of existing utility functions and interface to the existing data base. On this basis tools with enhanced functionality are provided. This framework has become standard for data analyses of the previously and currently collected data.

Distributed Computing Systems and Experiences / 147

HEP Applications Experience with the European DataGrid Middleware and Testbed

A. De Salvo¹; A. Fanfani²; A. Tsaregorodtsev³; C. Grandi⁴; D. Boutigny⁵; D. Colling⁶; D. Groep⁷; E. Luppi⁸; F. Brochu⁹; F. Carminati¹⁰; F. Harris¹¹; G. Negri¹²; I. Augustin¹⁰; I. Stokes-Rees¹³; J.J. Blaising⁵; J. Closier¹⁰; K. Bos⁷

; L. Perini¹²; M. Reale¹; O. Maroney¹⁴; O. Smirnova¹⁵; P. Capiluppi²; P. Cerello¹⁶; R. Barbera¹⁷; S. Bagnasco¹⁶; S. Burke¹⁸; S. Resconi¹²; V. GARONNE¹⁹; W. van Leeuwen⁷

¹ INFN Rome 1

² INFN Bologna

³ Marseille

⁴ INFN-Bologna

⁵ LAPP Annecy

⁶ Imperial

⁷ NIKHEF

⁸ INFN Sezione di Ferrara

⁹ Cambridge

¹⁰ CERN

¹¹ CERN/Oxford

¹² INFN Milan

¹³ Oxford

¹⁴ Bristol

¹⁵ Lund

¹⁶ INFN Turin

¹⁷ INFN Catania

¹⁸ Rutherford Appleton Laboratory

¹⁹ CNRS-IN2P3-CPPM, MARSEILLE

Auteur(s) contact: s.burke@rl.ac.uk

The European DataGrid (EDG) project ran from 2001 to 2004, with the aim of producing middleware which could form the basis of a production Grid, and of running a testbed to demonstrate the middleware. HEP experiments (initially the four LHC experiments and subsequently BaBar and D0) were involved from the start in specifying requirements, and subsequently in evaluating the performance of the middleware, both with generic tests and through increasingly complex data challenges. A lot of experience has therefore been gained which may be valuable to future Grid projects, in particular LCG and EGEE which are using a substantial amount of the middleware developed in EDG. We report our experiences with job submission, data management and mass storage, information and monitoring systems, Virtual Organisation management and Grid operations, and compare them with some typical Use Cases defined in the context of LCG. We also describe some of the main lessons learnt from the project, in particular in relation to configuration, fault-tolerance, interoperability and scalability, as well as the software development process itself, and point out some areas where further work is needed. We also make some comments on how these issues are being addressed in LCG and EGEE.

Distributed Computing Systems and Experiences / 228

HEP@HOME - A distributed computing system based on BOINC

A. Amorim¹; J. Villate²; P. Andrade³

¹ FCUL Portugal / CERN

² FEUP Portugal

³ University of Porto

Project SETI@HOME has proven to be one of the biggest successes of distributed computing during the last years. With a quite simple approach SETI manages to process huge amounts of data using a vast amount of distributed computer power.

To extend the generic usage of these kinds of distributed computing tools, BOINC (Berkeley Open Infrastructure for Network Computing) is being developed. In this communication we propose a BOINC version tailored to the specific requirements of the High Energy Physics (HEP) community - the HEP@HOME

The HEP@HOME will be able to process large amounts of data using virtually unlimited computing power, as BOINC does, and it should be able to work according to HEP specifications.

One of the main applications of distributed computing is distributed data analysis. In HEP the amounts of data to be analyzed have a large order of magnitude. Therefore, one of the design principles of this tool is to avoid data transfer - computation is done where data is stored. This will allow scientists to run their analysis applications even if they do not have a local copy of the data to be analyzed, taking advantage of either very large farms of dedicated computers or using their colleagues desktop PCs. This tool also satisfies other important requirements in HEP, namely, security, fault-tolerance and monitoring.

Wide Area Networking / 301

HEPBook - A Personal Collaborative HEP notebook

G. Roediger¹ ; P. POMATTO²

¹ CORPORATE COMPUTER SERVICES INC. - FERMILAB

² FNAL, Batavia, USA

Auteur(s) contact: roediger@fnal.gov

A High Energy Physics experiment has between 200 and 1000 collaborating physicists from nations spanning the entire globe. Each collaborator brings a unique combination of interests, and each has to search through the same huge heap of messages, research results, and other communication to find what is useful.

Too much scientific information is as useless as too little. It is time consuming, tedious, and difficult to sift and search for the pertinent bits. Often, the exact words to search for are unknown, or the information is badly organized, and the pertinent bits are not found. The search is abandoned, the time is lost, and valuable information is never communicated as it was intended.

Much of a collaboration's information is in the individual physicists paper logbooks. The physicists record important and pertinent information for their research. They save the log books to refer to it later, copy pages, and distribute them to their collaborators who share their interest and research.

Electronic Logbooks are now used in the control room of large detectors during the acquisition phase. They have proven useful for communicating the status of the detector and to keep the history of lab sessions in a format that can be queried and retrieved quickly. It has enabled remote monitoring of the detector and remote emergency help.

We have implemented an electronic Control Room Logbook, called CRL. It is used in the D0 experiment's detector control room for the Run II acquisition. As of mid

April 2004 there are over 305,000 entries in the D0 logbook, all viewable and able to be annotated from the web. Other experiments such as CMS, MiniBoone, and Minos have also adapted the CRL. These experiments all have very different needs, so they all configured and customized the CRL in many different ways. The HEPBook will move the logbook from the control room to the personal and collaborative HEP notebook. In this paper we will review the HEPBook technology and capabilities and discuss the new HEPBook architecture. Among the topics discussed will be the use of Java reflection to recursively produce an XML representation of an entry, the ability to save personal entries as well as share entries among a collaboration through multiple repositories which incorporate software agent technology, interface with the GRID, and implement multiple security models. The HEPBook runs on all Java platforms including Apple, Win32, and Linux. A brief demo will be given of the HEPBook.

Poster Session 3 - Tableau: 37 / 448

Harp data and software migration from Objectivity to Oracle

A. Valassi¹ ; D Geppert¹ ; D Kolev² ; E. Tcherniaev¹ ; K. Nienartowicz¹ ; M Lubeck¹ ; M Nowak¹

¹ CERN

² University of Sofia

Auteur(s) contact: andrea.valassi@cern.ch

The migration of the Harp data and software from an Objectivity-based to an Oracle-based data storage solution is reviewed in this presentation.

The project, which was successfully completed in January 2004, involved three distinct phases. In the first phase, which profited significantly from the previous COMPASS data migration project, 30 TB of Harp raw event data were migrated in two weeks to a hybrid persistency solution, storing raw event records in standard "flat" files and the corresponding metadata in Oracle as relational tables.

In the second phase, the longest to achieve in spite of the relatively limited data volume to migrate, the complex data model of Harp event collections was reimplemented for the Oracle-based solution. The relational schema design and the implementation of read-only navigational access to event collections in the Harp software framework using Oracle are reviewed in detail in the presentation. The third phase was the easiest, as it involved the migration of conditions data (time-varying non-event data) from the Objectivity to the Oracle implementation of a same C++ API, which acted as a screening layer between the data model and its implementation.

Distributed Computing Services / 500

Housing Metadata for the Common Physicist Using a Relational Database

A. Baranovski¹ ; A. Kreymer² ; A. Lyon² ; A. Sill³ ; C. Cioffi⁴ ; C. Nicholson⁵ ; F. Ratnikov⁶ ; G. Garzoglio² ; G. McCance⁷ ; I. Terekhov¹ ; J. Fulachier⁸ ; J. Trumbo⁹ ; K. Nienartowicz⁷ ; L. Loebel Carpenter¹ ; L. Lueking² ; M. Burgon-Lyon¹⁰ ; M. Leslie⁴ ; M. Zimmer¹ ; P. Vokac¹ ; P. Kunszt⁷ ; P. Millar⁵ ; R. Herber² ; R. Illingworth² ; R. Kennedy² ; R. Rocha⁷ ; R. St.Denis¹¹ ; S. Albrand¹² ; S. Belforte¹³ ; S. Hanlon¹⁰ ; S. Stonjek¹ ; S. Veseli¹ ; S. White¹ ; T. Barrass¹⁴ ; T. Doyle¹⁵ ; U. Kerzel¹⁶ ; V. Bartsch⁴ ; W. Merritt²

¹ *Fermilab*² *FERMI NATIONAL ACCELERATOR LABORATORY*³ *Texas Tech University*⁴ *Oxford University*⁵ *UNIVERSITY OF GLASGOW*⁶ *Rutgers University*⁷ *CERN*⁸ *LPSC, Grenoble*⁹ *FERMILAB*¹⁰ *Glasgow University*¹¹ *University of Glasgow*¹² *Grenoble*¹³ *INFN/Trieste*¹⁴ *University of Bristol*¹⁵ *Glasgow University*¹⁶ *Karlsruhe University***Auteur(s) contact:** wyatt@fnal.gov

SAM was developed as a data handling system for Run II at Fermilab. SAM is a collection of services, each described by metadata. The metadata are modeled on a relational database, and implemented in ORACLE. SAM, originally deployed in production for the D0 Run II experiment, has now been also deployed at CDF and is being commissioned at MINOS. This illustrates that the metadata decomposition of its services has a broader applicability than just one experiment. A joint working group on metadata with representatives from ATLAS, BaBar, CDF, CMS, D0, and LHCb in cooperation with EGEE has examined this metadata decomposition in the light of general HEP user requirements.

Greater understanding of the required services of a performant data handling system has emerged from Run II experience. This experience is being merged with the understanding being developed in the course of LHC experience with data challenges and user case discussions. We describe the SAM schema and the commonalities of function and service support between this schema and proposals for the LHC experiments. We describe the support structure required for SAM schema updates, the use of development, integration, and production instances. We are also looking at the LHC proposals for the evolution of schema using keyword-value pairs that are then transformed into a normalized, performant database schema.

Event Processing / 172

How to build an event store - the new Kanga Event Store for BaBar

M. Steinke¹¹ *Ruhr Universitaet Bochum*

In the past year, BaBar has shifted from using Objectivity to using ROOT I/O as the basis for our primary event store. This shift required a total reworking of Kanga, our ROOT-based data storage format. We took advantage of this opportunity to ease the use of the data by supporting multiple access modes that make use of many of the analysis tools available in ROOT.

Specifically, our new event store supports: 1) the pre-existing separated transient + persistent model, 2) a transient based load-on-demand model

currently being developed, 3) direct access to persistent data classes in compiled code, 4) fully interactive access to persistent data classes from either the ROOT prompt and via interpreted macros.

We will describe key features of Kanga including: 1) the separation and management of transient and persistent representations of data, 2) the implementation of read on demand references in ROOT, 3) the modular and extensible persistent event design, 4) the implementation of schema evolution and 5) BaBar specific extensions to core ROOT classes that we used to preserve the end-user “feel” of ROOT.

Distributed Computing Services / 536

Huge Memory systems for data-intensive science

R. MOUNT¹

¹ SLAC

Event Processing / 57

IGUANA Interactive Graphics Project: Recent Developments

G. Alverson¹ ; G. Eulisse¹ ; I. Osborne² ; L. Taylor¹ ; L. Tuura¹ ; S. MUZAFFAR³

¹ NORTHEASTERN UNIVERSITY

² Northeastern University, Boston, USA

³ NorthEastern University, Boston, USA

Auteur(s) contact: ianna.osborne@cern.ch

This paper describes recent developments in the IGUANA (Interactive Graphics for User ANALysis) project. IGUANA is a generic framework and toolkit, used by CMS and D0, to build a variety of interactive applications such as detector and event visualisation and interactive GEANT3 and GEANT4 browsers.

IGUANA is a freely available toolkit based on open-source components including Qt, OpenInventor (Coin3D) and OpenGL and LCG services.

New features we describe since the last CHEP conference include:

multi-document architecture;

user interface to Python scripting;

2D visualisation with auto-generation of slices/projections from 3D data;

per-object actions such as clipping, slicing, lighting or animation;

correlated actions (e.g. picking) for multiple views;

production of high-quality and compact vector postscript output from any OpenGL display, with surface shading and invisible surface culling (together with the gl2ps project).

We compare the IGUANA rendering, memory performance, and porting issues for various platforms including: Linux on x86, Windows, and Mac OSX with its native Quartz-Extreme rendering system.

Poster Session 3 - Tableau: 6 / 282

INDICO - the software behind CHEP 2004

H. Sanchez¹ ; J.B. Gonzalez Lopez¹ ; J-Y. Le Meur¹ ; T. Baron¹ ; V Turney¹

¹ CERN

Auteur(s) contact: thomas.baron@cern.ch, jean-yves.le.meur@cern.ch

CHEP 2004 conference is using the Integrated Digital Conferencing product to manage part of its web site and processes to run the conference.

This software has been built in the framework of InDiCo European Project. It is designed to be generic and extensible with the goal of providing help for single seminars as well as large conferences management. Partly developed at CERN within the Document Server (CDS) team, it focuses on supporting future events in HEP domain and it will be distributed with an open source license.

The presentation will explain the main application features before going into the details of its object oriented development. It will be demoed how the InDiCo can be used as a platform for scheduling events in a large institution like CERN and how it can give to conference chairperson a solid basis to set up, run and archive content of meetings.

INTAS / 534

INTAS Discussion

INTAS^{None}

INTAS (<http://www.intas.be>): International Association for the promotion of co-operation with scientists from the New Independent States of the former Soviet Union (NIS). INTAS encourages joint activities between its INTAS Members and the NIS in all exact and natural sciences, economics, human and social sciences.

INTAS supports a number of NIS participants to attend the 2004 Computing in High Energy Physics Conference (CHEP'04). During CHEP'04, this discussion has been organised so that NIS delegates can meet specifically with their physicists and computer scientists counterparts to discuss topics of mutual interest.

Provisional agenda

- Welcome message from Wolfgang Von Rueden
- Present status about East-West science cooperation inside the High Energy Physics Environment
- Main directions to go
- New developments and technologies needed to achieve the goals, especially in the network infrastructure area
- Time and financial schedules
- Needed man power
- Financial scenarios including funding scenarios (INTAS will need this as they plan for their future activities to focus on selected thematic)

IceTray: a Software Framework for IceCube

T. DeYoung¹

¹ UNIVERSITY OF MARYLAND

Auteur(s) contact: deyoung@pa.msu.edu

IceCube is a cubic kilometer-scale neutrino telescope under construction at the South Pole. The minimalistic nature of the instrument poses several challenges for the software framework. Events occur at random times, and frequently overlap, requiring some modifications of the standard event-based processing paradigm. Computational requirements related to modeling the detector medium necessitate the ability for software components to defer processing events. With minimal information from the detector, events must be reconstructed many times with different hypotheses or methods, and the results compared. The appropriate series of software components required to process an event varies considerably, and can be determined only at run time. Finally, reconstruction algorithms are constantly evolving, with development taking place throughout the collaboration, so it is essential that conversion of private analysis code to online production software be simple and, given the inaccessibility of the experimental site, robust. The IceCube collaboration has developed the IceTray framework, which meets these needs by blending aspects of push- and pull-based architectures to produce a highly modular system which nevertheless allows each software component a significant degree of control over the execution flow.

Core Software / 63

IgProf profiling tool

G. Eulisse¹ ; L. Tuura¹

¹ NORTHEASTERN UNIVERSITY OF BOSTON (MA) U.S.A.

Auteur(s) contact: giulio.eulisse@cern.ch

A fundamental part of software development is to detect and analyse weak spots of the programs to guide optimisation efforts. We present a brief overview and usage experience on some of the most valuable open-source tools such as valgrind and oprofile. We describe their main strengths and weaknesses as experienced by the CMS experiment.

As we have found that these tools do not satisfy all our needs, CMS has also developed a tool of its own called “igprof”. It complements the other tools, allowing us to profile memory usage, CPU usage, memory leaks and file descriptor usage of large complex applications such as the CMS reconstruction and analysis software. It is requires no instrumentation and works with multi-threaded programs and with all shared libraries, including dynamically loaded ones.

We describe this new tool, it’s features and output, and experience including improvements gained in CMS.

Event Processing / 117

Implementation and Performance of the High-Level Trigger electron and photon selection for the ATLAS experiment at the LHC

M. DIAZ GOMEZ^{None}

The ATLAS experiment at the Large Hadron Collider (LHC) will face the challenge of efficiently selecting interesting candidate events in pp collisions at 14 TeV center-of-mass energy, whilst rejecting the enormous number of background events, stemming from an interaction rate of about 10^9 Hz. The Level-1 trigger will reduce the incoming rate to around $O(100)$ kHz. Subsequently, the High-Level Triggers (HLT), which are comprised of the second level trigger and the event filter, will need to reduce this rate further by a factor of $O(10^3)$. The HLT selection is software based and will be implemented on commercial CPUs using a common framework, which is based on the standard ATLAS object-oriented software architecture. In this talk an overview of the current implementation of the selection for electrons and photons in the trigger is given. The performance of this implementation has been evaluated using Monte Carlo simulations in terms of the efficiency for the signal channels, the rate expected for the selection, the data preparation times, and the algorithm execution times. Besides the efficiency and rate estimates, some physics examples will be discussed, showing that the triggers are well adapted for the physics programme envisaged at LHC. The electron/gamma trigger software has been also integrated in the ATLAS 2004 combined test-beam, to validate the chosen selection architecture in a real on-line environment.

Computer Fabrics / 201

Implementation of a reliable and expandable on-line storage for compute clusters

J. VanWezel¹

¹ FORSCHUNGSZENTRUM KARLSRUHE

Auteur(s) contact: jvw@iwr.fzk.de

The HEP experiments that use the regional center GridKa will handle large amounts of data. Traditional access methods via local disks or large network storage servers show limitations in size, throughput or data management flexibility.

High speed interconnects like Fibre Channel, iSCSI or Infiniband as well as parallel file systems are becoming increasingly important in large cluster installations to offer the scalable size and throughput needed for PetaByte storage. At the same time the reliable and proven NFS protocol allows local area storage access via traditional Ethernet very cost effectively.

The cluster at GridKa uses the General Parallel File System (GPFS) on a 20 node file server farm that connects to over 1000 FC disks via a Storage Area Network. The 130 TB on-line storage is distributed to the 390 node cluster via NFS. A load balancing system ensures an even load distribution and additionally allows for on-line file server exchange.

Discussed are the components of the storage area network, specific Linux tools, and the construction and optimisation of the cluster file system along with the RAID groups. A high availability is obtained and measurements prove high throughput under different conditions. The use of the file system administration and management possibilities is presented as is the implementation and effectiveness of the load balancing system.

Plenary / 324

Improving Standard C++ for the Physics Community

M. Paterno¹ ; W. Brown¹

¹ *FERMILAB*

Auteur(s) contact: paterno@fnal.gov, wb@fnal.gov

As Fermilab's representatives to the C++ standardization effort, we have been promoting directions of special interest to the physics community. We here report on selected recent developments toward the next revision of the C++ Standard. Topics will include standardization of random number and special function libraries, as well as core language issues promoting improved run-time performance.

The random number library provides an extensible framework for random number generators. It includes a handful of widely-used and high-quality random number engines, as well as some of the most widely-used random number distributions. The modular design makes it easy for users to add their own engines, and perhaps more importantly their own distributions, on an equal footing with those in the library.

The special functions library contains many of the commonly-used functions of mathematical physics. These include a variety of cylindrical and spherical Bessel functions, Legendre and associated Legendre functions, hypergeometric and confluent hypergeometric functions, among others.

We also report on an ongoing analysis, and proposal for core language additions, with the goal of improved run-time performance. Current compilers routinely perform inter-procedural flow analysis within a compilation unit. These additions would allow compilers to perform comparable analysis between compilation units, and to optimize code based on their findings.

Poster Session 1 - Tableau: 32 / 152

InGRID - Installing GRID

F.M. Taurino¹ ; G. Tortone² ; R. Esposito²

¹ *INFN - INFN*

² *INFN*

Auteur(s) contact: taurino@na.infn.it

The "gridification" of a computing farm is usually a complex and time consuming task. Operating system installation, grid specific software, configuration files customization can turn into a large problem for site managers.

This poster introduces InGRID, a solution used to install and maintain grid software on small/medium size computing farms.

Grid elements installation with InGRID consists in three steps.

In the first step nodes are installed using RedHat Kickstart, an installation method that automate most of a Linux distribution installation, including disk partitioning, boot loader configuration, network configuration, base package selection.

Grid specific software is then integrated using apt4rpm, a package management wrapper over the rpm commands. Apt automatically manages packages dependencies, and is able to download, install and upgrade RPMs from a central software repository.

Finally, grid configuration files are customized through LCFGng, a system to setup and maintain Unix machines, that can configure many system files, execute scripts, create users, etc.

Computer Fabrics / 487

InfiniBand for High Energy Physics

A. Heiss¹ ; G. GANIS² ; U. Schwickerath³

¹ FORSCHUNGSZENTRUM KARLSRUHE

² CERN

³ Forschungszentrum Karlsruhe

Auteur(s) contact: andreas.heiss@iwr.fzk.de

Distributed physics analysis techniques as provided by the rootd and proofd concepts require a fast and efficient interconnect between the nodes. Apart from the required bandwidth the latency of message transfers is important, in particular in environments with many nodes. Ethernet is known to have large latencies, between 30 and 60 micro seconds for the common Giga-bit Ethernet.

The InfiniBand architecture is a relatively new, open industry standard. It defines a switched high-speed, low-latency fabric designed to connect compute nodes

and I/O nodes with copper or fibre cables. The theoretical bandwidth is up to 30 Gbit/s. The Institute for Scientific Computing (IWR) at the Forschungszentrum Karlsruhe is testing InfiniBand technology since begin of 2003, and has a cluster of

dual Xeon nodes using the 4X (10 Gbit/s) version of the interconnect.

Bringing the RFIO protocol - which is part of the CERN CASTOR facilities for sequential file transfers - to InfiniBand has been a big success, allowing significant reduction of CPU consumption and increase of file transfer speed.

A first prototype of a direct interface to InfiniBand for the root toolkit has been designed and implemented.

Experiences with hard- and software, in particular MPI performance results, will be reported.

The methods and first performance results on rfio and root will be shown and compared to other fabric technologies like Ethernet.

Distributed Computing Services / 417

Information and Monitoring Services within a Grid Environment

A. Cooke¹ ; A. Djaoui² ; A. Gray¹ ; A. Wilson² ; B. Coghlan³ ; D. O'Callaghan³ ; J. Leake⁴ ; J. Magowan⁵ ; J. Walk² ; L. Cornwall² ; M. Craig² ; N. Podhorszki⁶ ; O. Lyttleton³ ; P. Taylor⁵ ; R. Byrom² ; R. Cordenonsi⁷ ; R. Middleton² ; S. Fisher⁸ ; S. Hicks² ; S. Kenny³ ; W. Nutt¹

¹ Heriot Watt

² Rutherford Appleton Laboratory

³ Trinity College Dublin

⁴ Objective Engineering Ltd.

⁵ IBM

⁶ SZTAKI

⁷ Queen Mary, University of London

⁸ RAL

Auteur(s) contact: s.m.fisher@rl.ac.uk

The R-GMA (Relational Grid Monitoring Architecture) was developed within the EU DataGrid project, to bring the power of SQL to an information and monitoring system for the grid. It provides producer and consumer services to both publish and retrieve information from anywhere within a grid environment. Users within a Virtual Organization may define their own tables dynamically into which to publish data.

Within the DataGrid project R-GMA was used for the information system, making details about grid resources available for use by other middleware components. R-GMA has also been used for monitoring grid jobs by members of the CMS and D0 collaborations where information about jobs is published from within a job wrapper, transported across the grid by R-GMA and made available to users. An accounting package for processing PBS logging data and sending it to one or more Grid Operation Centres using R-GMA has been written and is being deployed within LCG. There are many other existing and potential applications.

R-GMA is currently being re-engineered to fit into a Web Service environment as part of the EU EGEE project. Improvements being developed include fine grained authorization, an improved user interface and measures to ensure superior scaling behaviour.

Poster Session 2 - Tableau: 26 / 250

Installing and Operating a Grid Infrastructure at DESY

A. Campbell¹ ; A. Gellrich¹ ; B. Lewendel¹ ; C. Wissing² ; F. Brasolin³ ; J Ferrando⁴ ; K. Wrona¹ ; M. Ernst¹ ; M. Vorobiev⁵ ; M. de Riese¹ ; P. Fuhrmann¹ ; R. Mankel¹ ; S. Padhi¹ ; U. Ensslin¹ ; V. Gulzow¹

¹ DESY

² Dortmund University

³ INFN, Bologna

⁴ Glasgow University

⁵ ITEP Moscow

Auteur(s) contact: andreas.gellrich@desy.de

DESY is one of the world-wide leading centers for research with particle accelerators and a center for research with synchrotron light. The hadron-electron collider HERA houses four experiments which are taking data and will be operated until 2006 at least.

The computer center manages a data volumes of order 1 PB and is the home for around 1000 CPUs.

In 2003 DESY started to set up a Grid infrastructure on site. Monte Carlo production is the primer HEP application candidate for the Grid at DESY. The experiments have started major tests.

A first Grid Testbed was based on EDG 1.4. Some effort was taken to install the binary distribution of the middleware on SuSE based Linux systems at DESY. With the first fixed LCG-2 release in spring 2004, the Grid Testbed2 was installed, which serves as the basis for all further DESY activities.

The contribution to CHEP2004 will start by briefly summarizing the status of the Grid activities at DESY in the context of EGEE and D-GRID, in which DESY takes a leading role.

In the following, we will discuss the integration of Grid components in the infrastructure of the DESY computer center.

This includes technical aspects of the operating system, such as SuSE versus RedHat Linux, the interaction with the mass storage system, and the management of Virtual Organizations.

We will finish with discussing installation and operation experiences of Grid middleware at DESY, also having in mind HEP and future synchrotron light experiments in the X-FEL era.

Poster Session 1 - Tableau: 30 / 221

Integrating Mutiple PC Farms into an uniform computing System with Maui

G. Sun¹

¹ *INSTITUE OF HIGH ENERGY PHYSICS*

Auteur(s) contact: sungx@mail.ihep.ac.cn

These are several on-going experiments at IHEP, such as BES, YBJ, and CMS collaboration with CERN. each experiment has its own computing system, these computing systems run separately. This leads to a very low CPU utilization due to different usage period of each experiment. The Grid technology is a very good candidate for integrating these separate computing systems into a "single image", but it is too early to be put into a production system as it is not stable and user-friendly as well. A realistic choice is to implement such an integration and sharing with Maui, an advacned scheduler. Each PC farm is thought as a partition, which is assigned high priority to its owner users with preemtor feature. this paper will describe the detail of implementation with Maui scheduler, as well as the entire system architecture and configuration and fuctions.

Online Computing / 331

Integration of ATLAS Software in the Combined Beam Test

M. Dobson¹

¹ *CERN*

Auteur(s) contact: marc.dobson@cern.ch

The ATLAS collaboration had a Combined Beam Test from May until October 2004. Collection and analysis of data required integration of several software systems that are developed as prototypes for the ATLAS experiment, due to start in 2007. Eleven different detector technologies were integrated with the Data Acquisition system and were taking data synchronously. The DAQ was integrated with the High Level Trigger software, which will perform online selection of ATLAS events. The data quality was monitored at various stages of the Trigger and DAQ chain. The data was stored in a format foreseen for ATLAS and was analyzed using a prototype of the experiments' offline software, using the Athena framework. Parameters recorded by the Detector Control System

were recorded in a prototype of the ATLAS Conditions Data Base and were made available for the offline analysis of the collected event data. The combined beam test provided a unique opportunity to integrate and to test the prototype of ATLAS online and offline software in its complete functionality.

Distributed Computing Systems and Experiences / 401

Interactive Data Analysis on the Grid using Globus 3 and JAS3

B. Ananthan¹ ; D. Alexander¹ ; M. Turri² ; T. Johnson² ; V. Serbo²

¹ *Tech-X corporation*

² *SLAC*

The aim of the service is to allow fully distributed analysis of large volumes of data while maintaining true (sub-second) interactivity. All the Grid related components are based on OGSA style Grid services, and to the maximum extent uses existing Globus Toolkit 3.0 (GT3) services. All transactions are authenticated and authorized using GSI (Grid Security Infrastructure) mechanism - part of GT3. JAS3, and experiment independent data analysis tool is used as the interactive analysis client.

The system consists of three main service components:

Dataset Catalog Service:

The Dataset Catalog supports browsing for an interesting dataset, or searching for data using a query language which operates on metadata stored in the catalog. The catalog makes few assumptions about the metadata stored in the catalog, except that the metadata consists of key-value pairs, stored in a hierarchical tree. The Dataset Catalog Service is designed to allow easy interfacing to existing data catalog back-ends.

Dataset Analysis Grid Service:

This service is responsible for resolving the dataset id from the catalog service, and transferring chunks of data to worker nodes for analysis processing. This service also manages the worker nodes, distributes analysis code to the worker nodes and retrieves intermediate results from the worker nodes before sending merged results back to the analysis client.

Worker Execution Services:

This service runs on each worker node and is responsible for processing analysis requests.

In this presentation we will demonstrate the current system, and will describe some of the choices made in architecting the system, in particular the challenges of obtaining interactive response times from GT3.

Event Processing / 255

Ion transport simulation using Geant4 hadronic physics

B. TRIEU¹ ; G. Folger² ; H.P. Wellisch² ; I. CORNELIU³ ; P. TRUSCOTT⁴ ; T. Koi⁵

¹ *CERN, Geneva, Switzerland*

² CERN

³ INFN, Torino, Italy

⁴ QinetiQ, Farnborough, UK

⁵ SLAC

Auteur(s) contact: tkoi@slac.stanford.edu, hpw@free.fr

The transportation of ions in matter is subject of much interest in not only high-energy ion-ion collider experiments such as RHIC and LHC but also many other field of science, engineering and medical applications. Geant4 is a tool kit for simulation of passage of particles through matter and its OO designs makes it easy to extend its capability for ion transports. To simulate ions interaction, we had to develop two major functionalities to Geant4. One is cross section calculators and the other is final stage generators for ion-ion interactions. For cross sections calculator, several empirical cross section formulas for the total reaction cross section of ion-ion interactions were investigated. And for final stage generator, binary cascade and quark-gluon string model of Geant4 were improved so that ions reaction with matter can also be calculated. Having successfully developed both functionalities, Geant4 can be applied to ion transportation problems. In the presentation we will explain cross section and final stage generator in detail and show comparisons with experimental data.

Event Processing / 399

JASSimApp plugin for JAS3: Interactive Geant4 GUI

E. Suzuki¹ ; H Yoshidah¹ ; K. Minamimoto¹ ; K. Murakami² ; M. Donszelmann³ ; M. Nagamatsu¹ ; M. Turri³ ; T. Johnson³ ; V. Serbo³

¹ *Naruto University of Education*

² *KEK*

³ *SLAC*

Auteur(s) contact: serbo@slac.stanford.edu, mark.donszelmann@slac.stanford.edu

JASSimApp is joint project of SLAC, KEK, and Naruto University to create integrated GUI for Geant4, based on JAS3 framework, with ability to interactively:

- Edit Geant4 geometry, materials, and physics processes
- Control Geant4 execution, local and remote: pass commands and receive output, control event loop
- Access AIDA histograms defined in Geant4
- Show generated Geant4 events in integrated event display

JAS3 is the latest development of JAS, a general-purpose data analysis tool. It employs a highly modular component-based framework and allows flexible and powerful customized plugin modules. JASSimApp is a concrete implementation of its design concept, Geant4 as the problem domain. It is composed of existing interactive tools like : GAG, Gain, Momo, WIRED etc.. A new C++ class of the Geant4 interfaces category was developed to exploit multi-threaded control over Geant4 execution. The plugin modules of JAS3 reused existing classes with little modification.

References:

<http://erpc1.naruto-u.ac.jp/~geant4/index.html> ; <http://jas.freehep.org/jas3> ; <http://wired.freehep.org> ; <http://wwwasd.web.cern.ch/wwwasd/geant4/geant4.html>

Poster Session 2 - Tableau: 12 / 293

JIM Deployment for the CDF Experiment

A. Baranovski¹ ; A. Kreymer² ; A. Lyon² ; A. Sill³ ; F. Ratnikov⁴ ; G. Garzoglio² ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter¹ ; L. Lueking² ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber² ; R. Illingworth² ; R. Kennedy² ; R. St.Denis⁷ ; S. Belforte⁸ ; S. Stonjek⁷ ; S. Veseli¹ ; S. White¹ ; U. Kerzel⁹ ; V. Bartsch⁶ ; W. Merritt²

¹ *Fermilab*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

³ *Texas Tech. University*

⁴ *Rutgers University*

⁵ *UNIVERSITY OF GLASGOW*

⁶ *Oxford University*

⁷ *Glasgow University*

⁸ *INFN, Trieste*

⁹ *Karlsruhe University*

Auteur(s) contact: m.burgon-lyon@physics.gla.ac.uk

JIM (Job and Information Management) is a grid extension to the mature data handling system called SAM (Sequential Access via Metadata) used by the CDF, DZero and Minos Experiments based at Fermilab. JIM uses a thin client to allow job submissions from any computer with Internet access, provided the user has a valid certificate or kerberos ticket. On completion the job output can be downloaded using a web interface. The JIM execution site software can be installed on shared resources, such as ScotGRID, as it may be configured for any batch system and does not require exclusive control of the hardware. Resources that do not belong entirely to CDF and thus cannot run DCAF (Decentralised CDF Analysis Farm), may therefore be accessed using JIM. We will report on the initial deployment of JIM for CDF and the steps taken to integrate JIM with DCAF.

Online Computing / 157

Jefferson Lab Data Acquisition Run Control System

C. Timmer¹ ; D. Abbott¹ ; D. LAWRENCE² ; E. Wolin¹ ; Ed Jastrzembski¹ ; G. Heyes¹ ; V. Gyurjyan¹

¹ *Jefferson Lab*

² *TJNAF, USA*

Auteur(s) contact: gurjyan@jlab.org

A general overview of the Jefferson Lab data acquisition run control system is presented. This run control system is designed to operate the configuration, control, and monitoring of all Jefferson Lab experiments. It controls data-taking activities by coordinating the operation of DAQ sub-systems, online software components and third-party software such as external slow control systems. The main, unique feature which sets this system apart from conventional systems is its incorporation of intelligent agent concepts. Intelligent agents are autonomous programs which interact with each other through certain protocols on a peer-to-peer level. In this case, the protocols and standards used come from the domain-independent Foundation for Intelligent Physical Agents (FIPA), and the implementation used is the Java Agent Development Framework (JADE). A lightweight, RDF (Resource Definition Framework) based language was developed to standardize the description of the run control system for configuration purposes. Fault tolerance and recovery issues are addressed.

Key features of the system include: subsystem state management, configuration management, agent communication, multiple simultaneous run management and synchronization, and user interfaces. A user interface allowing web-wide monitoring was developed which incorporates a JAS/AIDA data server extensible through Java servlets.

Poster Session 2 - Tableau: 3 / 271

Job Interactivity using a Steering Service in an Interactive Grid Analysis Environment

A. Ali¹ ; A. Anjum¹ ; C. Steenberg² ; F. van Lingen² ; H. Newman³ ; I. Willers⁴ ; J. Bunn⁵ ; M. A. Zafar¹ ; M. Thomas² ; R. Cavanaugh⁶ ; R. McClatchey⁷

¹ NIIT

² CALIFORNIA INSTITUTE OF TECHNOLOGY

³ Caltech

⁴ CERN

⁵ CALTECH

⁶ University of Florida

⁷ UWE

Auteur(s) contact: ashiq.anjum@niit.edu.pk

In the context of Interactive Grid-Enabled Analysis Environment (GAE), physicists desire bi-directional interaction with the job they submitted. In one direction, monitoring information about the job and hence a “progress bar” should be provided to them. On other direction, physicist should be able to control their jobs. Before submission, they may direct the job to some specified resource or computing element. Before execution, its parameter may be changed or it may be moved to another location. During execution, its intermediate results should be fetched or it may be moved to another location. Also, physicists should be able to kill, restart, hold and resume their jobs.

Interactive job execution requires that at each step, the user must make choices between alternative application components, files, or locations. So a dead end may be reached where no solution can be found, which would require backtracking to undo some previous choice. Another desire is reliable and optimal execution of the job. Grid should take some decisions regarding the job execution to help in reliable and optimal execution of the job. Reliability can be achieved using the job recovery mechanism. When a job on grid fails, the recovery mechanism should resubmit the job on either the same resource or on different resource. Check-pointing the job will make resource utilization low when recovering the job from failure.

In this paper the architecture and design of an autonomous grid service is described that fulfills the above stated requirements for interactivity in Grid-enabled data analysis.

Poster Session 2 - Tableau: 1 / 272

Job Monitoring in Interactive Grid Analysis Environment

A. Ali¹ ; A. Anjum¹ ; C. Steenberg² ; F. van Lingen² ; H. Newman³ ; I. Willers⁴ ; J. Bunn³ ; M. Thomas³ ; R. Cavanaugh⁵ ; R. McClatchey⁶ ; W. Rehman¹

¹ NIIT

² CALIFORNIA INSTITUTE OF TECHNOLOGY

³ Caltech

⁴ CERN

⁵ University of South Florida

⁶ UWE

Auteur(s) contact: ashiq.anjum@niit.edu.pk

Grid is emerging as a great computational resource but its dynamic behaviour makes the Grid environment unpredictable. System failure or network failure can occur or the system performance can degrade. So once the job has been submitted monitoring becomes very essential for user to ensure that the job is completed in an efficient way. In current environments once user submits a job he loses direct control over the job, system behaves like a batch system, user submits the job and gets the result back. Only information a user can obtain about a job is whether it is scheduled, running, cancelled or finished. This information is enough from the Grid management point of view but not from the point of view of a user. User wants interactive environment in which he can check the progress of the job, obtain intermediate results, terminate the job based on the progress of job or intermediate results, steer the job other nodes to achieve better performance and check the resources consumed by the job. So a mechanism is needed that can provide user with secure access to information about different attributes of a job. In this paper we describe a monitoring service, a java based web service that will provide secure access to different attributes of a job once a job has been submitted to Interactive Grid Analysis Environment.

Distributed Computing Systems and Experiences / 86

Job Submission & Monitoring on the PHENIX Grid*

A. Shevel¹ ; A. Withers¹ ; B. Jacak¹ ; D. Morrison² ; I. Sourikova³ ; M. Reuter¹ ; R. Lacey¹ ; T. Thomas⁴

¹ STATE UNIVERSITY OF NEW YORK AT STONY BROOK

² Brookhaven National Laboratory

³ Brookhaven National Laboratory

⁴ State University of New Mexico

Auteur(s) contact: shevel@mail.chem.sunysb.edu

The PHENIX collaboration records large volumes of data for each experimental run (now about 1/4 PB/year). Efficient and timely analysis of this data can benefit from a framework for distributed analysis via a growing number of remote computing facilities in the collaboration. The grid architecture has been, or is being deployed at most of these facilities.

The experience being obtained in the transition to the Grid infrastructure with minimum of manpower is presented with particular emphasis on job monitoring and job submission in multi cluster environment. The integration of the existing subsystems (from Globus project, from several HEP collaborations), large application libraries, and other software tools to render the resulting architecture stable, robust, and useful for the end user is also discussed.

Poster Session 2 - Tableau: 20 / 396

Job-monitoring over the Grid with GridIce infrastructure.

A. Pierro¹ ; G. Donvito² ; G. Maggi¹ ; G. Tortone³ ; M. Maggi¹ ; N. De Filippis¹ ; T. Coviello¹

¹ *INFN Bari*

² *UNIVERSITÀ DEGLI STUDI DI BARI*

³ *INFN Napoli*

Auteur(s) contact: giacinto.donvito@ba.infn.it

In a wide-area distributed and heterogeneous grid environment, monitoring represents an important and crucial task. It includes system status checking, performance tuning, bottlenecks detecting, troubleshooting, fault notifying. In particular a good monitoring infrastructure must provide the information to track down the current status of a job in order to locate any problems. Job monitoring requires interoperation between the monitoring system and other grid services.

Currently development and deployment LCG testbeds integrate GridICE monitoring system which measures and publishes the state of a grid resource at a particular point in time. In this paper we present the efforts to integrate in the current GridICE infrastructure, additional useful information about job status, e.g. the name of job, the virtual organization to which it belongs, eventually real and mapped user who has submitted the job, the effective CPU time consumed and its exit status.

Poster Session 2 - Tableau: 21 / 456

K5 @ INFN.IT: an infrastructure for the INFN cross REALM & AFS cell authentication.

E.M.V. Fasanelli¹

¹ *INFN.*

Auteur(s) contact: enrico.m.v.fasanelli@le.infn.it

The infn.it AFS cell has been providing a useful single file-space and authentication mechanism for the whole INFN, but the lack of a distributed management system, has lead several INFN sections and LABs to setup local AFS cells. The hierarchical transitive cross-realm authentication introduced in the Kerberos 5 protocol and the new versions of the OpenAFS and MIT implementation of Kerberos 5, make possible to setup an AFS cross cell authentication in a transparent way, using the Kerberos 5 cross-realm one. The goal of the K5 @ INFN.IT project is to provide a Kerberos 5 authentication infrastructure for the INFN and cross-realm authentication to be used for the cross cell AFS authentication. In this work we describe the scenario, the results of various tests performed, the solution chosen and the status of the K5 @ INFN.IT project.

Core Software / 447

LCG Conditions Database Project Overview

A. Amorim¹ ; A. Valassi² ; D. Duellmann² ; D. Klose³ ; L. Pedro¹ ; N. Barros¹ ; S.A. Schmidt⁴ ; T. Franco¹ ; V. Tsulaia⁵

¹ *University of Lisboa*

² *CERN*

³ *Universidade de Lisboa, Portugal*

⁴ *University of Mainz*

⁵ *University of Pittsburg*

Auteur(s) contact: andrea.valassi@cern.ch

The Conditions Database project has been launched to implement a common persistency solution for experiment conditions data in the context of the LHC Computing Grid (LCG) Persistency Framework. Conditions data, such as calibration, alignment or slow control data, are non-event experiment data characterized by the fact that they vary in time and may have different versions. The LCG project draws on preexisting projects which have led to the definition of a generic C++ API for condition data access and its implementation using different storage technologies, such as Objectivity, MySQL or Oracle. The project is assigned the task to deliver a production release of the software including implementation libraries for several technologies and high level tools for data management. The presentation will review the current status of the LCG common project at the time of the conference and the plans for its evolution.

Event Processing / 409

LCG Generator

A. RIBON¹ ; A. SHERSTNEV² ; A. VOLOGDIN² ; F. AMBROGLINI³ ; F. MOORTGAT¹ ; H. NAVES⁴ ; J. CUEVAS MAESTRO⁴ ; L. DUDKO² ; M. KIRSANOV⁵ ; P. Bartalini⁶ ; S. BELOV⁷ ; S. KOROBOV⁷ ; S. MAKARYCHEV⁸ ; V. UZHINSKY⁷

¹ *CERN, Geneva, Switzerland*

² *SINP MSU, Moscow, Russia*

³ *Perugia University, Italy*

⁴ *Cantabria University, Santander, Spain*

⁵ *INR RAS Troitsk, Moscow, Russia*

⁶ *CERN*

⁷ *JINR, Dubna, Russia*

⁸ *ITEP, Moscow, Russia*

Auteur(s) contact: paolo.bartalini@cern.ch

In the framework of the LCG Simulation Project, we present the Generator Services Sub-project, launched in 2003 under the oversight of the LHC Monte Carlo steering group (MC4LHC). The goal of the Generator Services Subproject is to guarantee the physics generator support for the LHC experiments. Work is divided into four work packages: Generator library; Storage, event interfaces and particle services; Public event files and event database; Validation and tuning. The current status and the future plans in the four different work packages are presented. Some emphasis is put on the Monte

Carlo Generator Library (GENSER) and on the Monte Carlo Generator Database (MCDB).

GENSER is the central code repository for Monte Carlo generators and generator tools. It was the first CVS repository in the LCG Simulation project and it is currently distributed in AFS. GENSER comprises release and building tools for librarian and end users. GENSER is going to gradually replace the obsolete CERN library in Monte Carlo generators support.

MCDB is a public database for the configuration, book-keeping and storage of the generator level event files. The generator events often need to be prepared and documented by Monte Carlo experts. MCDB aims at facilitating the communication between Monte-Carlo experts and end-users. Its use can be optionally extended to the official event production of the LHC experiments.

Core Software / 248

LCIO persistency and data model for LC simulation and reconstruction

F. Gaede¹ ; H. Vogt² ; N. Graf³ ; R. Cassell³ ; T. Behnke² ; T. Johnson³

¹ *DESY IT*

² *DESY*

³ *SLAC*

Auteur(s) contact: frank-dieter.gaede@cern.ch

LCIO is a persistency framework and data model for the next linear collider. Its original implementation, as presented at CHEP 2003, was focused on simulation studies. Since then the data model has been extended to also incorporate prototype test beam data, reconstruction and analysis. The design of the interface has also been simplified. LCIO defines a common abstract user interface (API) in Java, C++ and Fortran in order to fulfill the needs of the global linear collider community. It is designed to be lightweight and flexible without introducing additional dependencies on other software packages.

User code is completely separated from the concrete persistency implementation. SIO, a simple binary format that supports data compression and pointer retrieval is the current choice. LCIO is implemented in such a way that it can also be used as the transient data model in any linear collider application, e.g. a modular reconstruction program

can use the LCIO event class (LCEvent) as the container for the modules' input and output data. As LCIO offers a common API for three languages it is also possible to construct a multi-language reconstruction framework that would facilitate the integration of already existing algorithms.

A number of groups has already incorporated LCIO in their software frameworks and others plan to do so.

We present the design and implementation of LCIO, focusing on new developments and uses.

Poster Session 2 - Tableau: 49 / 364

LEXOR, the LCG-2 Executor for the ATLAS DC2 Production System

A. De Salvo¹ ; C. Negri² ; D. Rebatto³ ; L. Vaccarossa⁴

¹ *INFN Roma*

² *CNAF Bologna*

³ *INFN - MILANO*

⁴ *INFN Milano*

Auteur(s) contact: david.rebatto@mi.infn.it

In this paper we present an overview of the implementation of the LCG interface for the ATLAS production system. In order to take profit of the features provided by DataGRID software, on which LCG is based, we implemented a Python module, seamless integrated into the Workload Management System, which can be used as an object-oriented API to the submission services. On top of it we implemented Lexor, an executor component conforming to the pull/push model designed by the DC2 production system team. It pulls job descriptions from the supervisor component and uses them to create job objects, which in turn are submitted to the Grid. All the typical Grid operations (match-making with respect to input data location, registration of output data in the replica catalog, workload balancing) are performed by the underlying middleware, while interactions with ATLAS metadata catalog and the production database are granted by the integration with the Data Management System (Don Quijote) client module and via XML messages to the production supervisor (Windmill).

Poster Session 2 - Tableau: 64 / 289

LHC data files meet mass storage and networks: going after the lost performance

L. Tuura¹ ; T. Barrass² ; V. Innocente³

¹ *NORTHEASTERN UNIVERSITY, BOSTON, MA, USA*

² *Bristol University, UK*

³ *CERN, PH/SFT*

Auteur(s) contact: lassi.tuura@cern.ch

Experiments frequently produce many small data files for reasons beyond their control, such as output splitting into physics data streams, parallel processing on large farms, database technology incapable of concurrent writes into a single file, and constraints from running farms reliably. Resulting data file size is often far from ideal for network transfer and mass storage performance. Provided that time to analysis does not significantly deteriorate, files arriving from a farm could easily be merged into larger logical chunks, for example by physics stream and file type within a configurable time and size window.

Uncompressed zip archives seem an attractive candidate for such file merging and are currently tested by the CMS experiment. We describe the main components now in use: the merging tools, tools to read and write zip files directly from C++, plug-ins to the database system, mass-storage access optimisation, consistent handling of application and replica metadata, and integration with catalogues and other grid tools. We report on the file size ratio obtained in the CMS 2004 data challenge and observations and analysis on

changes to data access as well as estimated impact on network usage.

Wide Area Networking / 359

LambdaStation: A forwarding and admission control service to interface production network facilities with advanced research network paths

D. Petravick¹ ; P. DeMar¹

¹ *FERMILAB*

Auteur(s) contact: demar@fnal.gov

Advanced optical-based networks have the capacity and capability to meet the extremely large data movement requirements of particle physics collaborations. To date, research efforts in the advanced network area have been primarily focused on provisioning, dynamically configuring, and monitoring the wide area optical network infrastructure itself. Application use of these facilities has been largely limited to demonstrations using prototype high performance computing systems. Fermilab has initiated a project to enable our production network facilities to exploit these advanced research networks. Our objective is to selectively forward designated data transfers, on a per-flow basis, between capacious production-use storage systems on local campus networks, using a dynamically provisioned alternate path on a wide area advanced research network. To accomplish this, it is necessary to develop the capability to dynamically reconfigure forwarding of specific flows within our local production-use routers, provide an interface that enables applications to utilize the service, and dynamically implement appropriate access control on the alternate network path. Our project involves developing that infrastructure. We call it LambdaStation. If one envisions wide area optical network paths as high bandwidth data railways, then LambdaStation would functionally be the railroad terminal that regulates which flows at the local site get directed onto the high bandwidth data railways. LambdaStation is in a very early stage of development. Our paper will discuss its design, early deployment experiences, and future directions for the project.

Computer Fabrics / 346

Lattice QCD Clusters at Fermilab

A. Singh¹ ; D. Holmgren² ; J. Simone¹ ; P. Mackenzie¹

¹ *Fermilab*

² *FERMILAB*

Auteur(s) contact: djholm@fnal.gov

As part of the DOE SciDAC “National Infrastructure for Lattice Gauge Computing” project, Fermilab builds and operates production clusters for lattice QCD simulations. We currently operate three clusters: a 128-node dual Xeon Myrinet cluster, a 128-node Pentium 4E Myrinet cluster, and a 32-node dual Xeon Infiniband cluster. We will discuss the operation of these systems and examine their performance in detail. We will describe the uniform user runtime environment emerging from the SciDAC collaboration.

The design of lattice QCD clusters requires careful attention towards balancing memory bandwidth, floating point throughput, and network performance. We will discuss our investigations of various commodity processors, including Pentium 4E, Xeon, Itanium2, Opteron, and PPC970, in terms of their suitability for building balanced QCD clusters. We will also discuss our early experiences with the emerging Infiniband and PCI Express architectures. Finally, we will examine historical trends in price to performance ratios of lattice QCD clusters, and we will present our predictions and plans for future clusters.

Distributed Computing Services / 196

Lattice QCD Data and Metadata Archives at Fermilab and the International Lattice Data Grid

E. Neilsen¹ ; J. Simone¹

¹ *FERMI NATIONAL ACCELERATOR LABORATORY*

Auteur(s) contact: neilsen@fnal.gov

The lattice gauge theory community produces large volumes of data. Because the data produced by completed computations form the basis for future work, the maintenance of archives of existing data and metadata describing the provenance, generation parameters, and derived characteristics of that data is essential not only as a reference, but also as a basis for future work. Development of these archives according to uniform standards both in the data and metadata formats provided and in the software interfaces to the component services could greatly simplify collaborations between institutions and enable the dissemination of meaningful results.

This paper describes the progress made in the development of a set of such archives at the Fermilab lattice QCD facility. We are coordinating the development of the interfaces to these facilities and the formats of the data and metadata they provide with the efforts of the international lattice data grid (ILDG) metadata and middleware working groups, whose goals are to develop standard formats for lattice QCD data and metadata and a uniform interface to archive facilities that store them. Services under development include those commonly associated with data grids: a service registry, a metadata database, a replica catalog, and an interface to a mass storage system. All services provide GSI authenticated web service interfaces following modern standards, including WSDL and SOAP, and accept and provide data and metadata following recent XML based formats proposed by the ILDG metadata working group.

Poster Session 1 - Tableau: 25 / 317

Linux for the CLEO-c Online system

H. Schwarthoff¹ ; T. Wilksen¹

¹ *CORNELL UNIVERSITY*

Auteur(s) contact: hubert@mail.lns.cornell.edu

The CLEO collaboration at the Cornell electron positron storage ring CESR has completed its transition to the CLEO-c experiment. This new program contains a wide array of Physics studies of e^+e^- collisions at center of mass energies between 3 GeV and 5 GeV.

New challenges await the CLEO-c Online computing system, as the trigger rates are expected to rise from < 100 Hz to around 300 Hz at the J/Psi production threshold, with a moderate increase in data throughput requirements. While the current Solaris and VxWorks based readout system will perform adequately under those conditions, there is a desire to improve the performance of the central components to extend monitoring capabilities and provide larger safety margins. The solution, as in most modern particle detector systems, is to deploy Linux on Intel architecture computers for the performance critical applications.

For reasons of hardware and software availability, the existing CLEO Online and Offline computing environment has been ported to the Linux platform. This development allows the described challenge to be met.

In this presentation, we will report on our experiences adapting the CLEO Online computing system for operation under Linux. Issues regarding third party software and code portability will be addressed. Performance measurements will be presented.

Poster Session 2 - Tableau: 35 / 90

MONARC2: A Processes Oriented, Discrete Event Simulation Framework for Modelling and Design of Large Scale Distributed Systems.

H. Newman¹ ; I. Legrand²

¹ Caltech

² CALTECH

Auteur(s) contact: iosif.legrand@cern.ch

The design and optimization of the Computing Models for the future LHC experiments, based on the Grid technologies, requires a realistic and effective modeling and simulation of the data access patterns, the data flow across the local and wide area networks, and the scheduling and workflow created by many concurrent, data intensive jobs on large scale distributed systems.

This paper presents the latest generation of the MONARC (MOdels of Networked Analysis at Regional Centers) simulation framework, as a design and modelling tool for large scale distributed systems applied to HEP experiments. A process-oriented approach for discrete event simulation is used for describing concurrent running programs, as well as the stochastic arrival patterns that characterize how such systems are used. The simulation engine is based on Threaded Objects, (or Active Objects) which offer great flexibility in simulating the complex behavior of distributed data processing programs. The engine provides an appropriate scheduling mechanism for the Active Objects with efficient support for interrupts.

The framework provides a complete set of basic components (processing nodes, data servers, network components) together with dynamically loadable decision units (scheduling or data replication modules) for easily building complex Computing Model simulations.

Examples of simulating complex data processing systems, specific for the LHC

experiments (production tasks associated with data replication and interactive analysis on distributed farms) are presented, and the way the framework is used to compare different decision making algorithms or to optimize the overall Grid architecture.

Distributed Computing Services / 190

Managed Data Storage and Data Access Services for Data Grids

D. Petravick¹ ; I. Fisk¹ ; J. Bakken¹ ; M. Ernst² ; P. Fuhrmann² ; T. Mkrtchyan² ; T. Perelmutov¹

¹ *FERMILAB*

² *DESY*

Auteur(s) contact: michael.ernst@desy.de, ian.fisk@cern.ch

The LHC needs to achieve reliable high performance access to vastly distributed storage resources across the network. USCMS has worked with Fermilab-CD and DESY-IT on a storage service that was deployed at several sites. It provides Grid access to heterogeneous mass storage systems and synchronization between them. It increases resiliency by insulating clients from storage and network failures, and facilitates file sharing and network traffic shaping.

This new storage service is implemented as a Grid Storage Element (SE). It consists of dCache as the core storage system and an implementation of the Storage Resource Manager (SRM), that together allow both local and Grid based access to the mass storage facilities. It provides advanced functionalities for managing, accessing and distributing collaboration data.

USCMS is using this system both as Disk Resource Manager at Tier-1 and Tier-2 sites, and as Hierarchical Resource Manager with Enstore as tape back-end at the Fermilab Tier-1. It is used for providing shared managed disk pools at sites and for streaming data between the CERN Tier-0, the Fermilab Tier-1 and U.S. Tier-2 centers

Applications can reserve space for a time period, ensuring space availability when the application runs. Worker nodes without WAN connection can trigger data replication to the SE and then access data via the LAN. Moving the SE functionality off the worker nodes reduces load and improves reliability of the compute farm elements significantly.

We describe architecture, components, and experience gained in CMS production and the DC04 Data Challenge.

Managing software licences for a large research laboratory

A. Thys¹ ; H. Hansbakk¹ ; M. Katzarova¹ ; N. Hoemyr¹ ; P. Baehler¹ ; P. Saiz¹ ; T. Osborne¹

¹ *CERN IT*

Auteur(s) contact: nils.hoimyr@cern.ch

The Product Support (PS) group of the IT department at CERN distributes and supports more than one hundred different software packages, ranging from tools for computer aided design, field calculations, mathematical and structural analysis to software development. Most of these tools, which are used on a variety of Unix and Windows platforms by different user populations, are commercial packages requiring a licence. The group is also charged with license negotiations with the software vendors.

Keeping track of large number and variety of licences is no easy task, so in order to provide a more automated and more efficient service, the PS group has developed a database system to both track detailed licence configurations and to monitor their use. The system is called PSLicmon (PS Licence Monitor) and is based on an earlier development from the former CE group.

PSLicmon consists of four main components: report generation, data loader, Oracle product database and a PHP-based Web-interface. The license log parser/loader is implemented in Perl and loads reports from the different license managers into the Oracle database. The database contains information about products, licenses and suppliers and is linked to CERN's human resource database. The web-interface allows for on the fly generation of statistics plots as well as data entry and updates. The system also includes an alarm system for licence expiry.

Thanks to PSLicmon, the support team is able to better match licence acquisitions with the diverse needs of its user community, and to be in control of migration and phaseout scenarios between different products and/or product versions. The tool has proved to be a useful aid when making decisions regarding product support policy and licence acquisitions, in particular ensuring the provision of the correct number of often expensive software licences to match CERN's needs.

Poster Session 3 - Tableau: 46 / 277

Managing third-party software for the LCG

Auteur(s): E. Poinsignon¹

Co-auteur(s) Andreas Pfeiffer¹ ; Ioannis Papadopoulos¹ ; Pere Mato Vila¹ ; Saul Youssef² ; Shaun Ashby¹ ; Yannick Patois³

¹ *CERN*

² *BU*

³ *EGEE*

Auteur(s) contact: eric.poinsignon@cern.ch

The External Software Service of the LCG SPI project provides open source and public domain packages required by the LCG projects and experiments. Presently, more than 50 libraries and tools are provided for a set of platforms decided by the architect forum. All packages are installed following a standard

procedure and are documented on the web.
A set of scripts has been developed to ease new installations.

In addition to providing these packages, a software configuration management “toolbox” is provided, containing a coherent set of package-version combinations for each release of a project, as well as a distribution script which manages the dependencies of the LCG projects such that users can easily download and install a release of a project including its depended packages. Emphasis here has been put on ease of use for the end-user.

Poster Session 3 - Tableau: 53 / 155

Mantis: the Geant4-based simulation specialization of the CMS COBRA framework

A. DE ROECK¹ ; M. Liendl² ; M. Stavriankou³ ; P. ARCE⁴ ; S. BANERJEE⁵ ; T. BOCCALI⁶ ; T. TODOROV⁷ ; V. Innocente⁸

¹ *CERN, Geneva, Switzerland*

² *CERN*

³ *FNAL*

⁴ *CIEMAT, Madrid, Spain*

⁵ *Tata Institute, Bombay, India*

⁶ *SNS, Pisa, Italy*

⁷ *IRES, Strasbourg, France*

⁸ *CERN, PH/SFT*

Auteur(s) contact: maya.stavriankou@cern.ch

The CMS Geant4-based Simulation Framework, Mantis, is a specialization of the COBRA framework, which implements the CMS OO architecture. Mantis, which is the basis for the CMS-specific simulation program OSCAR, provides the infrastructure for the selection, configuration and tuning of all essential simulation elements: geometry construction, sensitive detector and magnetic field management, event generation and Monte Carlo truth, physics, particle propagation and tracking, run and event management, and user monitoring actions.

The experimental setup is built by Mantis using the COBRA Detector Description Database, DDD, which allows transparent instantiation of any layout (full or partial CMS simulation, test beam setups etc).

Persistency, histogramming and other important services are available using the standard COBRA infrastructure and are transparent to user applications.

Authors:

M. Stavriankou, P. Arce, S. Banerjee, T. Boccali, A. De Roeck, V. Innocente, M. Liendl, T. Todorov

Poster Session 2 - Tableau: 62 / 208

Mass Storage Management and the Grid

A. Earl¹ ; P. Clark¹

¹ *University of Edinburgh*

Auteur(s) contact: aearl@ed.ac.uk

The University of Edinburgh has an significant interest in mass storage systems as it is one of the core groups tasked with the roll out of storage software for the UK's particle physics grid, GridPP. We present the results of a development project to provide software interfaces between the SDSC Storage Resource Broker, the EU DataGrid and the Storage Resource Manager. This project was undertaken in association with the eDikt group at the National eScience Centre, the Universities of Bristol and Glasgow, Rutherford Appleton Laboratory and the San Diego Supercomputing Center.

Poster Session 1 - Tableau: 3 / 370

Methodologies and techniques for analysis of network flow data

A. Bobyshev¹ ; M. Grigoriev²

¹ *FERMILAB*

² *Fermilab*

Auteur(s) contact: bobyshev@fnal.gov

Network flow data gathered on border routers and core network switch/routers is used at Fermilab for statistical analysis of traffic patterns, passive network monitoring, and estimation of network performance characteristics. Flow data is also a critical tool in the investigation of computer security incidents. Development and enhancement of flow-based tools is on-going effort. The current state of flow analysis is based on the open source Flow-Tools package. This paper describes the most recent developments in flow analysis at Fermilab. Our goal is to provide a multidimensional view of network traffic patterns, with a detailed breakdown based on site, experiment, domain, subnet, hosts, protocol, or application. The latest analysis tool provides a descriptive and graphical representation of network traffic broken down by combinations of experiment and DNS domain. The tool can be utilized in real-time mode, as well as to provide a historical view. Another tool analyzes flow data to provide performance characteristics of completed multistream GridFTP data transfers. The current prototype provides a web interface for dynamic administration of the flow reports. We will describe and discuss the new features that we plan on developing in future enhancements to our flow analysis tool set.

Distributed Computing Services / 247

Middleware for the next generation Grid infrastructure

A. Aimar¹ ; A. Di Meglio¹ ; A. Edlund² ; D. Groep³ ; E. Laure¹ ; F. Hemmer¹ ; F. Pacini⁴ ; F. Prell⁵ ; L. Guy¹ ; M. Barroso¹ ; M. Livny⁶ ; M. Sgaravatto⁵ ; O. Mulmo⁷ ; P. Buncic¹ ; P. Kunszt¹ ; S. Beco⁴ ; S. Fisher⁸

¹ *CERN*

² *KTH*

³ *NIKHEF*

⁴ *DATAMAT*

⁵ *INFN*

⁶ *UNIVERSITY OF WISCONSIN*

⁷ *KTH, Sweden*

⁸ *RAL*

Auteur(s) contact: erwin.laure@cern.ch

The aim of the EGEE (Enabling Grids for E-Science in Europe) is to create a reliable and dependable European Grid infrastructure for e-Science. The objective of the Middleware Re-engineering and Integration Research Activity is to provide robust middleware components, deployable on several platforms and operating systems, corresponding to the core Grid services for resource access, data management, information collection, authentication & authorization, resource matchmaking and brokering, and monitoring and accounting.

For achieving this objective, we developed an architecture and design of the next generation Grid middleware leveraging experiences and existing components mainly from AliEn, EDG, and VDT. The architecture follows the service breakdown developed by the LCG ARDA RTAG. Our goal is to do as little original development as possible but rather re-engineer and harden existing Grid services. The evolution of these middleware components towards a Service Oriented Architecture (SOA) adopting existing standards (and following emerging ones) as much as possible is another major goal of our activity.

A rapid prototyping approach has been adopted, providing a sequence of more sophisticated prototypes to the EGEE candidate applications coming from the LHC HEP experiments and the Biomedical field. The close feedback loop with applications via these prototypes is indispensable for achieving our ultimate goals of providing a reliable and dependable Grid infrastructure.

In this paper we will report on the architecture and design of the main Grid components and report on our experiences with early prototype systems.

Online Computing / 281

Migrating PHENIX databases from object to relational model

D. Morrison¹ ; I. Sourikova¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Auteur(s) contact: irina@bnl.gov

To benefit from substantial advancements in Open Source database technology and ease deployment and development concerns with Objectivity/DB, the Phenix experiment at RHIC is migrating its principal databases from Objectivity to a relational database management system (RDBMS). The challenge of designing a relational DB schema to store a wide variety of calibration classes was solved by using ROOT I/O and storing each calibration object opaquely as a BLOB (Binary Large Object). Calibration metadata is stored as built-in types to allow fast index-based database search. To avoid a database back-end dependency the application was made ODBC-compliant (Open DataBase Connectivity is a standard database interface). An existent well-designed calibration DB API allowed users to be shielded from the underlying database technology change. Design choices and experience with transferring a large amount of Objectivity data into relational DB will be presented.

Grid Security / 463**Mis-use Cases for the Grid**D. Skow¹¹ *FERMILAB***Auteur(s) contact:** dane@fnal.gov

There have been a number of efforts to develop use cases for the Grid to guide development and useability testing. This talk examines the value of “mis-use cases” for guiding the development of operational controls and error handling. A couple of the more common current network attack patterns will be extrapolated to a global Grid environment. The talk will walk through the various activities necessary for incident response and recovery and strive to be technology neutral.

Gedanken incident response exercises are being discussed among the HEP PKI infrastructure specialists, but a systems-wide approach to the issues and necessary tools is needed. Determining scope of incidents, performing forensics and containing the spread requires a much more distributed approach than our previous experiences. A new set of tools and communication patterns are likely to be needed. This talk will be aimed at applications and middleware developers as well as operations teams for grids.

As time allows, the talk will survey current grid testbed middleware, identify the current control points and responsibilities and suggest places where extensions or modifications would be beneficial.

Distributed Computing Services / 89**MonALISA: An Agent Based, Dynamic Service System to Monitor, Control and Optimize Grid based Applications.**C Grigoras¹ ; C. Cirstoiu¹ ; C. Dobre¹ ; H. Newman² ; I. Legrand³ ; M. Toarta¹ ; R. Voicu⁴¹ *Polytechnic University Bucharest*² *Caltech*³ *CALTECH*⁴ *CERN***Auteur(s) contact:** iosif.legrand@cern.ch

The MonALISA (MONitoring Agents in A Large Integrated Services Architecture) system is a scalable Dynamic Distributed Services Architecture which is based on the mobile code paradigm.

An essential part of managing a global system, like the Grids, is a monitoring system that is able to monitor and track the many site facilities, networks, and all the task in progress, in real time. MonALISA is designed to easily integrate existing monitoring tools and procedures and to provide this information in a dynamic, self describing way to any other services or clients.

The monitoring information gathered is essential for developing higher level services that provide decision support, and eventually some degree of automated decisions, to help maintain and optimize workflow through the Grid.

MonALISA is an ensemble of autonomous multi-threaded, agent-based subsystems which are registered as dynamic services and are able to collaborate and cooperate in performing a wide range of monitoring, data processing and control tasks in large scale distributed applications. We also present the development of specialized higher level services, implemented as distributed mobile agents in the MonALISA framework to control and globally optimize tasks as grid scheduling, real-time data streaming or effective file replication.

The system is currently used to monitor several large scale systems and provides detailed information for computing nodes, LAN and WAN network components, job execution and applications specific parameters. This distributed system proved to be reliable, able to correctly handle connectivity problems and is running around the clock on more than 120 sites.

Poster Session 2 - Tableau: 44 / 65

Monitoring CMS Tracker construction and data quality using a grid/web service based on a visualization tool

A. Regano¹ ; G. Zito² ; M.S. Mennea³

¹ *University Bari*

² *INFN BARI*

³ *UNIVERSITY & INFN BARI*

Auteur(s) contact: giuseppe.zito@ba.infn.it

The complexity of the CMS Tracker (more than 50 million channels to monitor) now in construction in ten laboratories worldwide with hundreds of interested people , will require new tools for monitoring both the hardware and the software. In our approach we use both visualization tools and Grid services to make this monitoring possible. The use of visualization enables us to represent in a single computer screen all those million channels at once. The Grid will make it possible to get enough data and computing power in order to check every channel and also to reach the experts everywhere in the world allowing the early discovery of problems .

We report here on a first prototype developed using the Grid environment already available now in CMS i.e. LCG2. This prototype consists on a Java client which implements the GUI for Tracker Visualization and a few data servers connected to the tracker construction database , to Grid catalogs of event datasets or directly to test beam setups data acquisition . All the communication between client and servers is done using data encoded in xml and standard Internet protocols.

We will report on the experience acquired developing this prototype and on possible future developments in the framework of an interactive Grid and a virtual counting room allowing complete detector control from everywhere in the world.

Poster Session 2 - Tableau: 15 / 471

Monitoring a Petabyte Scale Storage System

A. Moibenko¹ ; C-H. Huang¹ ; D. Petravick¹ ; E. Berman¹ ; J. Bakken¹ ; M. Zalokar¹

¹ *FERMILAB*

Auteur(s) contact: berman@fnal.gov

Fermilab operates a petabyte scale storage system, Enstore, which is the primary data store for experiments' large data sets. The Enstore system regularly transfers greater than 15 Terabytes of data each day. It is designed using a client-server architecture providing sufficient modularity to allow easy addition and replacement of hardware and software components. Monitoring of this system is essential to insure the integrity of the data that is stored in it and to maintain the high volume access that this system supports.

The monitoring of this distributed system is accomplished using a variety of tools and techniques that present information for use by a variety of roles (operator, storage system administrator, storage software developer, user).

All elements of the system are monitored: performance, hardware, firmware, software, network, data integrity.

We will present details of the deployed monitoring tools with an emphasis on the different techniques that have proved useful to each role. Experience with the monitoring tools and techniques, what worked and what did not will be presented.

Poster Session 1 - Tableau: 27 / 484

Monitoring the CDF distributed computing farms

A. Eskola¹ ; E. Lipeles² ; F. Wuerthwein³ ; I. Sfiligoi⁴ ; M. Neubauer⁵

¹ *Rovaniemi Polytechnic, Finland*

² *UNIVERSITY OF CALIFORNIA SAN DIEGO*

³ *University of California San Diego*

⁴ *INFN Frascati*

⁵ *University of California, San Diego*

Auteur(s) contact: lipeles@cern.ch

CDF is deploying a version of its analysis facility (CAF) at several globally distributed sites. On top of the hardware at each of these sites is either an FBSNG or Condor batch manager and a SAM data handling system which in some cases also makes use of dCache.

The jobs which run at these sites also make use of a central database located at Fermilab. Each of these systems has its own monitoring.

In order to maintain and effectively use the distributed system, it is important that both the administrators and the users can get a complete global view of the system.

We will present a system which integrates the monitoring of all of these services into one globally accessible system based on the Monalisa product. This system is intended for administrators to monitor the system status and service level and for users to better locate resources and monitor job progress.

In addition, it is meant to satisfy the request by the CDF International Finance Committee that global computing resource usage by CDF can be audited.

Poster Session 3 - Tableau: 18 / 406

Monte Carlo Event Generation in a Multilanguage, Multiplatform Environment

A. Johnson¹ ; N. Graf¹

¹ SLAC

Auteur(s) contact: norman.graf@slac.stanford.edu

We discuss techniques used to access legacy event generators from modern simulation environments. Examples will be given of our experience within the linear collider community accessing various FORTRAN-based generators from within a Java environment. Coding to a standard interface and use of shared object libraries enables runtime selection of generators, and allows for extension of the suite of available generators without having to rewrite core code.

Poster Session 2 - Tableau: 52 / 47

Multi-Terabyte EIDE Disk Arrays running Linux RAID5

D. Petravick¹ ; D. Sanders² ; D. Summers² ; L. Cremaldi² ; M.D. Joy² ; R. Godang² ; V. Eschenburg²

¹ Fermilab

² UNIVERSITY OF MISSISSIPPI

Auteur(s) contact: sanders@phy.olemiss.edu

High-energy physics experiments are currently recording large amounts of data and in a few years will be recording prodigious quantities of data. New methods must be developed to handle this data and make analysis at universities possible. Grid Computing is one method; however, the data must be cached at the various Grid nodes. We examine some storage techniques that exploit recent developments in commodity hardware. Disk arrays using RAID level 5 (RAID5) include both parity and striping. The striping improves access speed. The parity protects data in the event of a single disk failure, but not in the case of multiple disk failures. We report on tests of dual-processor Linux Software RAID5 arrays and Hardware RAID5 arrays using the 12-disk 3ware controller, in conjunction with 300 GB disks, for use in offline high-energy physics data analysis. The price of IDE disks is now less than \$1/GB. These RAID5 disk arrays can be scaled to sizes affordable to small institutions and used when fast random access at low cost is important.

Poster Session 3 - Tableau: 9 / 360

Muon Event Filter Software for the ATLAS Experiment at LHC

M. Biglietti¹

¹ UNIVERSITY OF MICHIGAN

Auteur(s) contact: michela.biglietti@cern.ch

At LHC the 40 MHz bunch crossing rate dictates a high selectivity of the ATLAS Trigger system, which has to keep the full physics potential of the experiment in spite of a limited storage capability.

The level-1 trigger, implemented in a custom hardware, will reduce the initial rate to 75 kHz and is followed by the software based level-2 and Event Filter, usually referred as High Level Triggers (HLT), which further reduce the rate to about 100 Hz.

In this paper an overview of the implementation of the offline muon reconstruction algorithms MOORE (Muon Object Oriented REconstruction) and MuId (Muon Identification) as Event Filter in the Atlas online framework is given.

The MOORE algorithm performs the reconstruction inside the Muon Spectrometer providing a precise measurement of the muon track parameters outside the calorimeters; MuId combines the measurements of all ATLAS sub-detectors in order to identify muons and provides the best estimate of their momentum at the production vertex.

In the HLT implementation the muon reconstruction can be executed in the “full scan mode”, performing pattern recognition in the whole muon spectrometer, or in the “seeded mode”, taking advantage of the results of the earlier trigger levels.

An estimate of the execution time will be presented along with the performances in terms of efficiency, momentum resolution and rejection power for muons coming from hadron decays and for fake muon tracks, due to accidental hit correlations in the high background environment of the experiment.

Event Processing / 227

Muon Reconstruction Software in CMS

N. Neumeister¹

¹ CERN / HEPHY VIENNA

Auteur(s) contact: norbert.neumeister@cern.ch

The CMS detector has a sophisticated four-station muon system made up of tracking chambers (Drift Tubes, Cathode Strip Chambers) and dedicated trigger chambers. A muon reconstruction software based on Kalman filter techniques has been developed which reconstructs muons in the standalone muon system, using information from all three types of muon detectors, and links the resulting muon tracks with tracks reconstructed in the silicon tracker. The software is designed to work for both, offline reconstruction and for online event selection within the CMS High-Level Trigger (HLT). Since the quality of the selection algorithms used in the HLT system is of utmost importance the software has been designed using modern object-oriented software techniques and is implemented within the CMS reconstruction software framework. The system should be able to select events with final-state muons, indicating interesting physics. The design implementation and performance of the CMS muon reconstruction software is presented. We will show that offline code with little modifications can be used in the HLT system, by making use of the concepts of regional and conditional reconstruction. The implementation and performance of possible HLT selection algorithms are illustrated.

Plenary / 144

Network Architecture: lessons from the past, vision for the future

F. Fluckiger¹

¹ CERN

Auteur(s) contact: francois.fluckiger@cern.ch

The Architectural Principles of the Internet have dominated the past decade. Orthogonal to the telecommunications industry principles, they dramatically changed the networking landscape because they relied on iconoclastic ideas. First, the Internet end-to-end principle, which stipulates that the network should intervene minimally on the end-to-end traffic, pushing the complexity to the end-systems. Second, the ban of centralized functions: all the Internet techniques (routing, DNS, management) are based on distributed, decentralized mechanisms. Third, the absolute domination of connectionless (stateless) protocols (as with IP, HTTP).

However, when facing new requirements: multimedia traffic, security, Grid applications, these principles appear sometimes as architectural barriers. Multimedia requires QoS guarantees, but stateless systems are not good at QoS. Security requires active, intelligent networks, but dumb routers or plain end-to-end mail systems are insufficient. Grid applications require middleware overlay networks, often with centralized functions.

Attempts to overcome these deficiencies may lead to excessively complicated hybrid solutions, distorting the initial principles (the QoS Pandora box). Middleware solutions are sometimes difficult to deploy (e.g for large scale PKI deployment). "Lambda on-demand" technologies are conceptually nothing else than old switched circuits, that we never managed to satisfactorily integrate with IP networks.

Where is all this going? To help forming a vision of the future, the paper will refer to several observations that the author has formulated over the past 30 years: the "breathing law" (a succession of decentralization and recentralization phases), the perpetual and oscillating mismatch of the bandwidth offer-demand, the conceptual antagonisms between resource level and complexity, between scaling and QoS.

Computer Fabrics / 465

Network Information and Management Infrastructure Project

D. LaMore¹ ; D. Petravick¹ ; D. Skow¹ ; E. Berman¹ ; I. Mandrichenko¹ ; M. Crawford¹ ; M. Grigoriev¹ ; P. DeMar¹ ; R. Reitz¹ ; V. Podstavkov¹

¹ FNAL

Auteur(s) contact: ivm@fnal.gov

Management of large site network such as FNAL LAN presents many technical and organizational challenges. This highly dynamic network consists of around 10 thousand network nodes. The nature of the activities FNAL is involved in and its computing policy require that the network remains as open as reasonably possible both in terms of connectivity to the outside networks and in with respect to procedural simplicity of joining the network by temporary network participants such as visitors notebook computers. The goal of the Network Information and Management Infrastructure

project at FNAL is to build software infrastructure which would help network management and computer security teams organize monitoring and management of the network, simplify communication between these entities and users, integrate network management into FNAL computer center management infrastructure.

Primary authors: Phil DeMar (FNAL), Igor Mandrichenko (FNAL), Don Petravick(FNAL), Dane Skow (FNAL)

Wide Area Networking / 121

Networking for High Energy and Nuclear Physics as Global E-Science

Auteur(s): H. Newman¹

Co-auteur(s) (Does Not Apply) on Behalf of the ICFA Standing Committee on Inter-Regional Connectivity²

¹ *Caltech*

² *ICFA*

Auteur(s) contact: harvey.newman@cern.ch

Wide area networks of sufficient, and rapidly increasing end-to-end capability are vital for every phase of high energy physicists' work. Our bandwidth usage, and the typical capacity of the major national backbones and intercontinental links used by our field have progressed by a factor of more than 1000 over the past decade, and the outlook is for a similar increase over the next decade, as we enter the era of LHC physics served by Grids on a global scale. Responding to these trends, and the emerging need to provide rapid access and distribution of Petabyte-scale datasets, physicists working with network engineers and computer scientists are learning to use networks effectively in the 1-10 Gigabit/range, placing them among the leading developers of global networks.

In this talk I review the network requirements and usage trends, and present a bandwidth roadmap for HEP and other fields of "data intensive" science. I give an overview of the status and outlook for the world's research networks, technology advances, and the problem of the Digital Divide, based on the recent work of ICFA's Standing Committee on Inter-regional Connectivity (SCIC). Finally, I discuss the role of high speed networks in the next generation of Grid systems that are now being constructed to support data analysis for the LHC experiments.

[This is a candidate Plenary Presentation.]

Poster Session 3 - Tableau: 51 / 265

New Applications of PAX in Physics Analyses at Hadron Colliders

A. Schmidt¹ ; C. Jung¹ ; C. Weiser¹ ; D. Hirschbuehl¹ ; G. Quast¹ ; J. Rehn¹ ; K. Rabbertz¹ ; M. Erdmann¹ ; M. Kirsch¹ ; P. Schemitz¹ ; S. Kappler¹ ; S. Schalla¹ ; T. Walter¹ ; U. FELZMANN²

¹ *Institut fuer Experimentelle Kernphysik, Karlsruhe University, Germany*

² *University of Karlsruhe, Germany*

Auteur(s) contact: alexander.schmidt@cern.ch

At CHEP03 we introduced “Physics Analysis eXpert” (PAX), a C++ toolkit for advanced physics analyses in High Energy Physics (HEP) experiments. PAX introduces a new level of abstraction beyond detector reconstruction and provides a general, persistent container model for HEP events. Physics objects like fourvectors, vertices and collisions can easily be stored, accessed and manipulated. Bookkeeping of relations between these objects (like decay trees, vertex and collision separation, including deep copies etc.) is fully provided by a “relation manager”. Event container and associated objects represent a uniform interface for algorithms and facilitate the parallel development and evaluation of different physics interpretations of individual events. So called “analysis factories”, which actively identify and distinguish different physics processes, can easily be constructed with the PAX toolkit.

PAX has been officially released to the experiments CDF (Tevatron) and CMS (LHC) during the last year. It is being explored by a growing user community and applied in various complex physics analyses. We report about the successful application in studies of $t\bar{t}$ production at the Tevatron and Higgs searches in the channel $t\bar{t}H$ at the LHC.

Poster Session 1 - Tableau: 14 / 211

New compact hierarchical mass storage system at Belle realizing a peta-scale system with inexpensive ice-raid disks and an S-ait tape library

J. Hikita¹ ; K. Goto¹ ; K. Tamura¹ ; M. Makino² ; M. Yokoyama³ ; N. Katayama⁴ ; T. Hibino³

¹ *SONY Co.*

² *Systemworks Co.*

³ *HIGH ENERGY ACCELERATOR RESEARCH ORGANIZATION*

⁴ *KEK*

The Belle experiment has accumulated an integrated luminosity of more than 240fb⁻¹ so far, and a daily logged luminosity now exceeds 800pb⁻¹. These numbers correspond to more than 1PB of raw and processed data stored on tape and an accumulation of the raw data at the rate of 1TB/day. To meet these storage demands, a new cost effective, compact hierarchical mass storage system has been constructed. The system consists of commodity RAID systems using IDE disks and Linux PC servers as the front-end and a tape library system using the new high density SONY S-AIT tape as the back-end. The SONY Peta Serv software manages migration and restoration of the files between tapes and disks. The capacity of the tape library is, at the moment, 500 TB in three 19 inch racks and the RAID system, 64 TB in two 19 inch racks. An extension of the system to 1.2 PB tape library in eight racks with 150 TB RAID in four racks is planned. In this talk, experiences with the new system will be discussed and the performance of the system when used for data processing and physics analysis of the Belle experiment will be demonstrated.

Poster Session 2 - Tableau: 4 / 91

New distributed offline processing scheme at Belle

F. Ronga¹ ; I. Adachi¹ ; N. Katayama¹

¹ KEK

Auteur(s) contact: ichiro.adachi@kek.jp

The Belle experiment has accumulated an integrated luminosity of more than 240fb⁻¹ so far, and a daily logged luminosity has exceeded 800pb⁻¹. This requires more efficient and reliable way of event processing. To meet this requirement, new offline processing scheme has been constructed, based upon technique employed for the Belle online reconstruction farm. Event processing is performed at PC farms, which consists of 60 quad(0.7GHz) and 225 dual(1.3GHz or 3.2GHz) CPU PC nodes. Raw event data are read from a Solaris tape server connected to a DTF2 tape drive, and they are distributed over all PC nodes. Reconstructed events are recorded onto 8 file servers, which are newly installed last year. To maximize processing capabilities, various optimizations such as PC clustering, job control, output data management and so on have been done. As a result, processing power with this scheme has been more than doubled, which corresponds to that more than 3 fb⁻¹ of beam data per day can be processed. In this talk, stable operation of our new system, together with a description of the Belle offline computing model, will be demonstrated by showing computing performance obtained from experience in processing beam data.

Online Computing / 252

New experiences with the ALICE High Level Trigger Data Transport Framework

H. Tilsner¹ ; T.M. Steinbeck¹ ; V. Lindenstruth¹

¹ KIRCHHOFF INSTITUTE OF PHYSICS, RUPRECHT-KARLS-UNIVERSITY HEIDELBERG, for the Alice Collaboration

Auteur(s) contact: timm.steinbeck@kip.uni-heidelberg.de

The Alice High Level Trigger (HLT) is foreseen to consist of a cluster of 400 to 500 dual SMP PCs at the start-up of the experiment. It's input data rate can be up to 25GB/s. This has to be reduced to at most 1.2 GB/s before the data is sent to DAQ through event selection, filtering, and data compression. For these processing purposes, the data is passed through the cluster in several stages and groups for successive merging until, at the last stage, fully processed complete events are available. For the transport of the data through the stages of the cluster, a software framework is being developed consisting of multiple components. These components can be connected via a common interface to form complex configurations that define the data flow in the cluster. For the framework, new benchmark results are available as well as experience from tests and data challenges run in Heidelberg.

The framework is scheduled to be used during upcoming testbeam experiments.

Poster Session 3 - Tableau: 32 / 297

New specific solids definitions in the Geant4 geometry modeller

O. Link¹

¹ CERN, PH/SFT

Auteur(s) contact: oliver.link@cern.ch

Twisted trapezoids are important components in the LAr end cap calorimeter of the ATLAS detector. A similar solid, the so-called twisted tube consists of two end planes, inner and outer hyperboloidal surfaces, and twisted surfaces, and is an indispensable component for cylindrical drift chambers (see K. Hoshina et al, Computer Physics Communications 153 (2003) 373-391). In Geant3 exists a general version of a twisted trapezoid, however the implementation puts very strong restrictions on its use.

In the Geant4 toolkit no solids have been available to date to describe twisted objects. The design and realisation of new twisted solids within the framework of Geant4 will be presented together with the algorithmic details, followed by a discussion of the performance and accuracy test results.

Event Processing / 275

Offline Software for the ATLAS Combined Test Beam

A. FARILLA¹ ; M. GALLAS²

¹ I.N.F.N. ROMA3

² CERN

Auteur(s) contact: ada.farilla@roma3.infn.it

A full slice of the barrel detector of the ATLAS experiment at the LHC is being tested this year with beams of pions, muons, electrons and photons in the energy range 1-300 GeV in the H8 area of the CERN SPS. It is a challenging exercise since, for the first time, the complete software suite developed for the full ATLAS experiment has been extended for use with real detector data, including simulation, reconstruction, online and offline conditions databases, detector and physics monitoring, and distributed analysis.

Important integration issues like combined simulation, combined reconstruction, connection with the online services and management of many different types of conditions data are being addressed for the first time, with the goal of both achieving experience on such integration aspects and of performing physics studies requiring the combined analysis of simultaneous data coming from different subdetectors. It is a unique opportunity to test, with real data, new algorithms for pattern recognition, particle tracking and identification and High Level Trigger strategies.

A relevant outcome of this combined test beam will be a detailed comparison of Monte Carlo - based on Geant4 - with real data. In the talk the main components of the

software suite are described, together with some preliminary results obtained both with simulated and real data.

Distributed Computing Services / 218

On Distributed Database Deployment for the LHC Experiments

A. Vaniachine¹ ; D. Duellmann² ; D. Malon¹ ; I. Bird² ; J. Shiers² ; M. Girone² ; M. Lamanna²

¹ ANL

² CERN

Auteur(s) contact: dirk.duellmann@cern.ch, jamie.shiers@cern.ch

While there are differences among the LHC experiments in their views of the role of databases and their deployment, there is relatively widespread agreement on a number of principles:

1. Physics codes will need access to database-resident data. The need for database access is not confined to middleware and services: physics-related data will reside in databases.
2. Database-resident data will be distributed, and replicated. A single, centralized database, at CERN or elsewhere, does not suffice.
3. Distributed deployment infrastructure should be open to the use of different technologies as appropriate at the various Tier N sites.

A variety of approaches to distributed deployment have been explored in the context of individual experiments; indeed, a degree of distributed deployment has been integral to the computing model tests of some experiments (cf. ATLAS) in their 2004 data challenges. Approaches to replication have also been investigated in the context of specific databases, often with vendor-specific replication tools (e.g., Oracle Replication via Streams for the LCG File Catalog and the Oracle instantiation of the LCG conditions database; MySQL tools for replication in the MySQL instantiation of the LCG conditions database). XML exchange mechanisms have also been discussed. Distributed database deployment, though, is more than a middleware and applications software issue—a successful strategy must involve those who will be responsible for systems deployment and administration at LHC grid sites.

We describe the status of ongoing work in this area, and discuss the prospects for components of a common approach to distributed deployment in the time frame of the 2005 LHC data challenges.

Poster Session 2 - Tableau: 70 / 363

On the Management of Certification Authority in Large Scale GRID Infrastructure

E. Berdnikov¹

¹ INSTITUTE FOR HIGH ENERGY PHYSICS, PROTVINO, RUSSIA

Auteur(s) contact: evgueni.berdnikov@cern.ch

The scope of this work is the study of scalability limits of the Certification Authority (CA), running for large scale GRID environments.

The operation of Certification Authority is analyzed from the view of the rate of incoming requests, complexity of authentication procedures, LCG security restrictions and other limiting factors. It is shown, that standard CA operational model has some native “bottlenecks”, which can be resolved with proper management and technical tools.

The central point is the discussion of “decentralized” scheme with single CA and multiple authentication agents, called Registration Authorities (RA). Single CA retains a role for technical center, responsible for support of GRID security infrastructure, while general role of RAs is verification of requests from end-users.

Practical implementation of this scheme (including the development and installation of end-user software) have been done in CERN in 2002 (<http://service-grid-ca.web.cern.ch/service-grid-ca/help/RA.html>). Second implementation of the same ideas was the GRID project of the Russia Ministry of Atomic Energy, 2003 (<http://grid.ihep.su/MAG/>). These two implementations are compared in aspects of security and functionality.

Wide Area Networking / 284

On-demand Layer VPN Support for Grid Applications

E. Ronchieri¹ ; S. Andreozzi¹ ; T. Ferrari¹

¹ INFN CNAF

Auteur(s) contact: elisabetta.ronchieri@cnafe.infn.it, tiziana.ferrari@cnafe.infn.it

The problem of finding the best match between jobs and computing resources is critical for an efficient work load distribution in Grids. Very often jobs are preferably run on the Computing Elements (CEs) that can retrieve a copy of the input files from a local Storage Element (SE). This requires that multiple file copies are generated and managed by a data replication system.

We propose the use of scheduled on-demand Layer 2 Virtual Private Networks (L2 VPNs) for an alternative data access model based on the possibility to connect to the same virtual LAN both CEs and SEs from remote Grid domains. The L2 VPN members are “close” to each other. In this way a CE can be selected by a Resource Broker without requiring the presence of a local file replica. This simplifies the data management and allows a more efficient use of the network resources on the links connecting the Grid to its main data sources.

In this paper we detail how L2 VPNs are dynamically provisioned through the Grid Network Agreement Service. We propose a hierarchical network resource abstraction, the Path, and we show how it can be integrated in the Grid Information Service to perform network resource discovery and matchmaking. We then describe the User Interface through which the Path negotiation terms are specified by the user and we propose a path management approach that integrates different network technologies, namely MPLS and the Differentiated Services architecture. The implementation details of a system

prototype are described together with some initial experimental results.

Poster Session 1 - Tableau: 20 / 466

Online Monitoring and online calibration/reconstruction for the PHENIX experiment

C. Pinkenburg¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Auteur(s) contact: pinkenburg@bnl.gov

The PHENIX experiment consists of many different detectors and detector types, each one with its own needs concerning the monitoring of the data quality and the calibration. To ease the task for the shift crew to monitor the performance and status of each subsystem in PHENIX we developed a general client server based framework which delivers events at a rate in excess of 100Hz.

This model was chosen to minimize the possibility of accidental interference with the monitoring tasks themselves. The user only interacts with the client which can be restarted any time without loss or alteration of information on the server side. It also enables multiple people to check simultaneously the same detector - if need be even from remote locations. The information is transferred in form of histograms which are processed by the client. These histograms are saved for each run and some html output is generated which is used later on to remove problematic runs from the offline analysis. An additional interface to a data base is provide to enable the display of long term trends.

This framework was augmented to perform an immediate calibration pass and a quick reconstruction of rare signals in the counting house. This is achieved by filtering out interesting triggers and processing them on a local Linux cluster. That enabled PHENIX to e.g. keep track of the number of J/Psi's which could be expected while still taking data.

Poster Session 3 - Tableau: 7 / 213

OpenPAW

G B. Barrand¹

¹ *CNRS / IN2P3 / LAL*

Auteur(s) contact: barrand@lal.in2p3.fr

OpenPAW is for people that definitively do not want to quit the PAW command prompt, but seek anyway an implementation based over more modern technologies. We shall present the OpenScientist/Lab/opaw program that offers a PAW command prompt by using the OpenScientist tools (then C++, Inventor for doing graphic, Rio for doing

the IO, OnX for the GUI, etc...). The OpenScientist/Lab package being also AIDA compliant, we shall show that it is possible to marry AIDA with a PAW command prompt.

Core Software / 311

OpenScientist. Status of the project.

G B. Barrand¹

¹ *CNRS / IN2P3 / LAL*

Auteur(s) contact: barrand@lal.in2p3.fr

We want to present the status of this project.

After quickly remembering the basic choices around GUI, visualization and scripting we would like to develop what had been done in order to have an AIDA-3.2.1 compliant system, to visualize Geant4 data (G4Lab module), to visualize ROOT data (Mangrove module), to have an hippodraw module and what had been done in order to run on MacOSX by using the native NextStep (Cocoa) environment.

Plenary / 186

Operating the LCG and EGEE Production Grids for HEP

I. Bird¹

¹ *CERN*

Auteur(s) contact: ian.bird@cern.ch

In September 2003 the first LCG-1 service was put into production at most of the large Tier 1 sites and was quickly expanded up to 30 Tier 1 and Tier 2 sites by the end of the year. Several software upgrades were made and the LCG-2 service was put into production in time for the experiment data challenges that began in February 2004 and continued for several months. In particular LCG-2 introduced transparent access to mass storage and managed disk-only storage elements, and a first release of the Grid File Access library. Much valuable experience was gained during the data challenges in all aspects from the functionality and use of the middleware, to the deployment, maintenance, and operation of the services at many sites. Based on this experience a program of work to address the functional and operational issues is being implemented. The goal is to focus on essential areas such as data management and to build by the end of 2004 a basic grid system capable of handling the basic needs of LHC computing, providing direction for future middleware and service development.

The LCG-2 infrastructure also forms the production service of EGEE. This involves supporting new application communities, bringing in new sites not associated with HEP and evolving a full scale 24x7 user and operational support structure. We will describe the EGEE infrastructure, how it supports and interacts with LCG, and how we expect the infrastructure to evolve over the next year of the EGEE project.

Computer Fabrics / 374

Operation of the CERN Managed Storage environment; current status and future directions

C. Curran¹ ; D. Hughes¹ ; G. Lee¹ ; H. Cacote¹ ; J. van Eldik¹ ; T. Osborne¹ ; T. Smith¹ ; V. Bahyl¹

¹ CERN

Auteur(s) contact: tim.smith@cern.ch

This paper discusses the challenges in maintaining a stable Managed Storage Service for users built upon dynamic underlying disk and tape layers.

Early in 2004 the tools and techniques used to manage disk, tape, and stage servers were refreshed in adopting the QUATTOR tool set. This has markedly increased the coherency and efficiency of the configuration of data servers. The LEMON monitoring suite was deployed to raise alarms and gather performance metrics. Exploiting this foundation, higher level service displays are being added, giving comprehensive and near-real-time views of operations. The scope of our monitoring has been broadened to include low-level machine sensors such as thermometer, IPMI and SMART readings, improving our ability to detect impending hardware failure.

In terms of operations, widespread disk reliability problems which were manpower intensive to chase, were overcome by exchanging a bad batch of 1200 disks. Recent LHC data challenges have ventured into new operating domains for the CASTOR system, with massive disk resident file catalogues requiring special handling. The tape layer has focused on STK 9940 drives for bulk recording capacity: a large scale data migration to this media permitted old drive technologies to be retired. Repacking 9940A data to 9940B high density media allows us to recycle tapes, giving substantial savings by avoiding acquisition of new media.

In addition to more robust software, hardware developments are required for LHC era services. We are moving from EIDE to SATA based disk storage and envisage a tape drive technology refresh. Details will be provided of our investigations in these areas.

Core Software / 160

Optimizing Selection Performance on Scientific Data by utilizing Bitmap Indices

H. Schmuecker¹

¹ CERN

Auteur(s) contact: helmut.gerhard.schmucker@cern.ch

Bitmap indices have gained wide acceptance in data warehouse applications handling large amounts of read only data. High dimensional ad hoc queries can be efficiently performed by utilizing bitmap indices, especially if the queries cover only a subset of the attributes stored in the database. Such access patterns are common use in HEP analysis. Bitmap indices have been implemented by several commercial database management systems. However, the provided query algorithms focus on typical business applications, which are based on discrete attributes with low cardinality. HEP data, which are mostly characterized by non discrete attributes, cannot be queried efficiently by these implementations.

Support for selections on continuously distributed data can be added to the bitmap index technique by extending it with an adaptive

binning mechanism. Following this approach a prototype has been implemented, which provides the infrastructure to perform index based selections on HEP analysis data stored in ROOT trees/tuples. For the indices a range encoded design with multiple components has been chosen. This design concept allows to realize a very fine binning granularity, which is crucial to selection performance, with an index of reasonable size. Systematic performance tests have shown that the query processing time and the disk-I/O can be significantly reduced compared to a conventional scan of the data. This especially applies to optimization scenarios in HEP analysis, where selections are slightly varied and performed repetitively on one and same data sample.

Poster Session 2 - Tableau: 43 / 56

OptorSim: a Simulation Tool for Scheduling and Replica Optimisation in Data Grids

C. Nicholson¹ ; D. Cameron¹ ; F. Zini² ; K. Stockinger³ ; P. Millar¹ ; R. Carvajal-Schiaffino²

¹ UNIVERSITY OF GLASGOW

² ITC-irst

³ Lawrence Berkeley National Laboratory

Auteur(s) contact: c.nicholson@physics.gla.ac.uk

In large-scale Grids, the replication of files to different sites is an important data management mechanism which can reduce access latencies and give improved usage of resources such as network bandwidth, storage and computing power.

In the search for an optimal data replication strategy, the Grid simulator OptorSim was developed as part of the European DataGrid project. Simulations of various HEP Grid scenarios have been undertaken using different job scheduling and file replication algorithms, with the experimental emphasis being on physics analysis use-cases. Previously, the CMS Data Challenge 2002 testbed and UK GridPP testbed were among those simulated; recently, our focus has been on the LCG testbed. A novel economy-based strategy has been investigated as well as more traditional methods, with the economic models showing distinct advantages in terms of improved resource usage.

Here, an overview of OptorSim's design and implementation is presented with a selection of recent results, showing its usefulness as a Grid simulator both in its current features and in the ease of extensibility to new scheduling and replication algorithms.

Event Processing / 320

Overview and new developments in Geant4 electromagnetic physics

H. BURKHARDT¹ ; L. Urban² ; M. Maire³ ; P. GUMPLINGER⁴ ; R.P. KOKOULIN⁵ ; V. Grichine² ; V. Ivantchenko⁶

¹ CERN, Geneva, Switzerland

² CERN

³ LAPP

⁴ TRIUMF, Vancouver, Canada

⁵ *MEPhI, Moscow, Russia*

⁶ *CERN, ESA*

Auteur(s) contact: vladimir.ivantchenko@cern.ch

We will summarize the recent and current activities of the Geant4 working group responsible of the standard package of electromagnetic physics. The major recent activities include an design iteration in energy loss and multiple scattering domain providing “process versus models” approach, and development of the following physics models: multiple scattering, ultra relativistic muon physics, photo-absorption-ionisation model, ion ionisation, optical processes. An automatic acceptance suite of validation of physics is under development. Also we will comment on evolution of the concept of physics list.

Core Software / 381

POOL Development Status and Plans

A. Vaniachine¹ ; C. Cioffi² ; D. Duellmann³ ; D. Malon⁴ ; G. Govi⁵ ; H. Schmuecker⁵ ; I. Papadopoulos⁵ ; J. Moscicki⁵ ; K. Karr¹ ; M. Frank⁵ ; M. Girone⁵ ; R. Chytracsek⁵ ; T. Barrass⁶ ; W. Tannenbaum⁷ ; Z. Xie⁸

¹ *Argonne National Lab*

² *Oxford University*

³ *CERN IT/DB & LCG POOL PROJECT*

⁴ *ANL*

⁵ *CERN*

⁶ *Bristol University*

⁷ *Fermilab*

⁸ *Princeton University*

Auteur(s) contact: dirk.duellmann@cern.ch

The LCG POOL project is now entering the third year of active development. The basic functionality of the project is provided but some functional extensions will move into the POOL system this year. This presentation will give a summary of the main functionality provided by POOL, which used in physics productions today. We will then present the design and implementation of the main new interfaces and components planned such as the POOL RDBMS abstraction layer and the RDBMS based Storage Manager backend.

Core Software / 382

POOL Integration into three Experiment Software Frameworks

A. Schaffer¹ ; A. Vaniachine² ; C. Cioffi³ ; D. Duellmann⁴ ; D. Malon⁵ ; G. Govi⁴ ; H. Schmuecker⁴ ; I. Papadopoulos⁴ ; J. Moscicki⁴ ; K. Karr² ; M. Frank⁴ ; M. Girone⁴ ; P van Gemmeren⁶ ; P. MATO⁴ ; R. Chytracsek⁴ ; T. Barrass⁷ ; V. Innocente⁸ ; W. Tannenbaum⁹ ; Z. Xie¹⁰

¹ *Orsay*

² *Argonne National Lab*

³ *Oxford University*

⁴ *CERN*

⁵ *ANL*

⁶ *Brookhaven National Lab*

⁷ *Bristol University*

⁸ *CERN, PH/SFT*

⁹ *Fermilab*

¹⁰ *Princeton University*

Auteur(s) contact: dirk.duellmann@cern.ch

The POOL software package has been successfully integrated with the three large experiment software frameworks of ATLAS, CMS and LHCb. This presentation will summarise the experience gained during these integration efforts and will try to highlight the commonalities and the main differences between the integration approaches. In particular we'll discuss the role of the POOL object cache, the choice of the main storage technology in ROOT (tree or named objects) and approaches to collection and catalogue integration.

Event Processing / 135

Panoramix

G B. Barrand¹

¹ *CNRS / IN2P3 / LAL*

Auteur(s) contact: barrand@lal.in2p3.fr

Panoramix is an event display for LHCb. LaJoconde is an interactive environment over DaVinci, the analysis software layer for LHCb. We shall present global technological choices behind these two softwares : GUI, graphic, scripting, plotting. We shall present the connection to the framework (Gaudi), how we can integrate other tools like hippodraw. We shall present the overall capabilities to these systems and their today status. We shall outline also how good part of choosen technologies may be reused to build the same kind of interactive environments for ATLAS (LAL Agora prototype).

Poster Session 3 - Tableau: 50 / 288

Parallel compilation of CMS software

G. Eulisse¹ ; L. Tuura¹ ; S. Ashby² ; S. Schmid³

¹ *Northeastern University, Boston, MA, USA*

² *CERN*

³ *ETH Zurich*

Auteur(s) contact: lassi.tuura@cern.ch

LHC experiments have large amounts of software to build. CMS has studied ways to shorten project build times using parallel and distributed builds as well as improved ways to decide what to rebuild. We have experimented with making idle desktop and server machines easily available as a virtual build cluster using distcc and zeroconf. We have also tested variations of ccache and more traditional make dependency analysis.

We report on our test results, with analysis of the factors that most improve or limit build performance.

Poster Session 1 - Tableau: 18 / 203

Parallel implementation of Parton String Model event generator

A. Asryan¹ ; A. Zarochensev² ; G. Feofilov³ ; P. Naumenko¹ ; S. Nemnyugin⁴ ; V. Solodkov⁵ ; V. Vechernin³ ; V. Zolotarev³

¹ *Postgraduate*

² *Researcher*

³ *Professor*

⁴ *ASSOCIATE PROFESSOR*

⁵ *Student*

Auteur(s) contact: nemnugin@mph.phys.spbu.ru

We report the results of parallelization and tests of the Parton String Model event generator at the parallel cluster of St.Petersburg State University Telecommunication center.

Two schemes of parallelization were studied. In the first approach master process coordinates work of slave processes, gathers and analyzes data. Results of MC calculations are saved in local files. Local files are sent to the host computer on which the program of data processing is started. The second approach uses the parallel write in the common file shared between all processes. In this case the load of a communication subsystem of the cluster grows. Both approaches are realized with MPICH library. Some problems including the pseudorandom number generation in parallel computations were solved.

The modified parallel version of the PSM code includes a number of the additional possibilities: a selection of the impact parameter windows, the account of acceptance of the experimental setup and trigger selection data, and the calculation of various long range correlations between such observables as mean transverse momentum and charged particles multiplicity.

Poster Session 2 - Tableau: 54 / 438

Participation of Russian sites in the Data Challenge of ALICE experiment in 2004

A. Selivanov¹ ; A. Zarochencev² ; E. Ryabinkin³ ; G. Shabratova⁴ ; V. Dobretsov³ ; V. Kolosov⁵ ; V. Korenkov⁴ ; V. Mitcin⁴ ; V. Zolotarev² ; Y. Bugaenko⁴ ; Y. Kharlov⁶ ; Y. Lyublev⁷

¹ ITEP, Moscow

² SPbSU

³ KIAE, Moscow

⁴ Joint Institute for Nuclear Research (JINR)

⁵ ITEP

⁶ State Res.Center of Russian Feder. HEP Inst (IFVE)

⁷ Institute for Theoretical and Experimental Physics (ITEP)

The report presents an analysis of the Alice Data Challenge 2004. This Data Challenge has been performed on two different distributed computing environments. The first one is the Alice Environment for distributed computing (AliEn) used standalone. Presently this environment allows ALICE physicists to obtain results on simulation, reconstruction and analysis of data in ESD format for AA and pp collisions at LHC energies. The second environment is the LCG-2 middleware accessed via AliEn with the help of an interface, developed at INFN. Three Russian sites have been configured as AliEn nodes for the Data Challenge. These sites (IHEP at Protvino, ITEP in Moscow and JINR at Dubna) could run a maximal of 86 jobs. The initial analysis shows that the architecture of one site was not adequate for distributed computing. Another farm had nodes with insufficient RAM for efficient job processing. All these problems have been cured subsequent DC phases.

Actions have also been taken to reduce the downtime due to wrong site configuration. The local AliEn server installed at the JINR site has been used as a standard configuration for the other Russian sites. The total number of jobs processed in Russia constitute ~2% of total run in the ALICE DC 2004.

Poster Session 1 - Tableau: 23 / 475

Patching PCs

C. Sieh¹ ; D. Skow¹ ; J. Schmidt¹ ; M. Crawford¹ ; T. Dawson¹ ; T. Doody¹

¹ Fermilab

FNAL has over 5000 PCs running either Linux or Windows software. Protecting these systems efficiently against the latest vulnerabilities that arise has prompted FNAL to take a more central approach to patching systems. We outline the lab support structure for each OS and how we have provided a central solution that works within existing support boundaries. The paper will cover how we identify what patches are considered crucial for a system on the FNAL network and how we verify that systems are appropriately patched.

Poster Session 3 - Tableau: 24 / 200

Paths: Specifying Multiple Job Outputs via Filter Expressions

C. Jones¹

¹ CORNELL UNIVERSITY

A common task for a reconstruction/analysis system is to be able to output different sets of events to different permanent data stores (e.g. files). This allows multiple related logical jobs to be grouped into one process and run using the same input data (read from a permanent data store and/or created from an algorithm). In our system, physicists can specify multiple output 'paths', where each path contains a group of filters followed by output 'operations'. The filters are combined using a physicist specified boolean expression; only if the expression evaluates to true will the output operation be performed for that event.

Paths do not explicitly contain the order that data objects should be created as our system uses a 'data on demand' mechanism which causes data to be created the first time the data is requested. Separating the data dependencies from the event selection criteria vastly simplifies the task of creating a path, thereby making the facility more accessible to physicists.

Poster Session 2 - Tableau: 42 / 341

Patriot: Physics Archives and Tools required to Investigate Our Theories

S. Mrenna¹

¹ *FERMILAB*

Auteur(s) contact: mrenna@fnal.gov

PATRIOT is a project that aims to provide better predictions of physics events for the high-Pt physics program of Run2 at the Tevatron collider.

Central to Patriot is an enstore or mass storage repository for files describing the high-Pt physics predictions. These are typically stored as StdHep files which can be handled by CDF and D0 and run through detector and triggering simulations. The definition of these datasets in the CDF and D0 data handling system SAM is under way.

Patriot relies heavily on a new generation of Monte Carlo tools (such as MadEvent, Alpgen, Grappa, CompHEP, etc.) to calculate the hard structure of high-Pt events and the more venerable event generators (Pythia and Herwig) to make particle level predictions.

An early informational database, describing the types of data files stored in Patriot, already exists. A new database is under development.

In parallel with PATRIOT, we wish to develop the QCD tools that describe the detailed properties of high-Pt events. Some of the essential features of particle-level events must be described by non-perturbative functions, whose form is often constrained by theory, but which must be ultimately tuned to data.

Computer Fabrics / 72

Performance analysis of Cluster File System on Linux

D. XU¹ ; G. CHEN¹ ; L. WANG¹ ; W. WU¹ ; Y. CHENG¹

¹ COMPUTING CENTER, INSTITUTE OF HIGH ENERGY PHYSICS, CHINESE ACADEMY OF SCIENCES

With the development of Linux and improvement of PC's performance, PC cluster used as high performance computing system is becoming much popular. The performance of I/O subsystem and cluster file system is critical to a high performance computing system. In this work the basic characteristics of cluster file systems and their performance are reviewed. The performance of four distributed cluster file systems, AFS, NFS, PVFS and CASTOR, were measured. The measurements were carried out on CERN version RedHat 7.3.3 Linux using standard I/O performance benchmarks. Measurements show that for single-server single client configuration, NFS, CASTOR and PVFS have better performance and write rate slightly increases while the record length becomes larger. CASTOR has the best throughput when the number of write processes increases. PVFS and CASTOR are tested on multi-server and multi-client system. The two file systems nicely distribute data I/O to all servers. CASTOR RPIO protocol shows the best utilization of network bandwidth and optimized to large data size files. CASTOR also has the better scalability as a cluster file system. Based on the test some methods are proposed to improve the performance of cluster file system.

Poster Session 2 - Tableau: 48 / 58

Performance of an operating High Energy Physics Data grid, D0SAR-grid

B. Abbott¹ ; B. Quinn² ; D. Meyer³ ; E. Gregores⁴ ; H. Kim³ ; H. Severini¹ ; J. Snow⁵ ; J. Yu⁶ ; M. Sosebee³ ; N. Mondal⁷ ; P. Baringer⁸ ; P. Skubic¹ ; S. Novaes⁴ ; T. Bolton⁹ ; Z. Greenwood¹⁰

¹ University of Oklahoma

² The University of Mississippi

³ University of Texas at Arlington

⁴ Universidade Estadual Paulista, Brazil

⁵ Langston University

⁶ UNIVERSITY OF TEXAS AT ARLINGTON

⁷ Tata Institute of Fundamental Research, India

⁸ University of Kansas

⁹ Kansas State University

¹⁰ Louisiana Tech. University

Auteur(s) contact: jaehoonyu@uta.edu

The D0 experiment at Fermilab's Tevatron will record several petabytes of data over the next five years in pursuing the goals of understanding nature and searching for the origin of mass. Computing resources required to analyze these data far exceed the capabilities of any one institution. Moreover, the widely scattered geographical distribution of collaborators poses further serious difficulties for optimal use of human and computing resources. These difficulties will be exacerbated in future high energy physics experiments, like those at the LHC. The computing grid has long been recognized as a solution to these problems. This technology is being made a more immediate reality to end users by developing a fully realized grid in the D0 Southern Analysis Region (D0SAR). D0SAR consists of eleven universities in the Southern US, Brazil, Mexico and India. The centerpiece of D0SAR is a data and resource hub, a Regional Analysis Center (RAC). Each D0SAR member institution constructs an Institutional Analysis Center (IAC), which acts as a gateway to the grid for users within that institution. These IACs combine dedicated rack-mounted servers and personal desktop computers into a local physics analysis cluster. D0SAR has been working on establishing an

operational regional grid, D0SAR-Grid, using all available resources within it and a home-grown local task manager, McFarm.

In this talk, we will describe the architecture of the D0SAR-Grid implementation, the use and functionality of the grid, and the experiences of operating the grid for simulation, reprocessing and analysis of data from a currently running HEP experiment.

Online Computing / 217

Performance of the ATLAS DAQ DataFlow system

A. CORSORADU¹ ; A. DOS ANJOS² ; A. KACZMARSKA³ ; A. KUGEL⁴ ; A. LANKFORD⁵ ; A. MISIEJUK⁶ ; B. DIGIROLAMO¹ ; B. GORINI¹ ; B. GREEN⁶ ; B. MARTIN¹ ; C. HAEBERLI⁷ ; C. HINKELBEIN⁴ ; C. MEIROSU¹ ; D. BOTTERILL⁸ ; D. FRANCIS¹ ; D. PRIGENT¹ ; E. PALENCIA CORTEZON¹ ; E. PASQUALUCCI⁹ ; F. WICKENS⁸ ; G. CRONE¹⁰ ; G. KIEFT¹¹ ; G. LEHMANN¹ ; G. MORNACCHI¹ ; G. unel¹² ; H. ZOBERNIG¹³ ; H.P. BECK⁷ ; J. DAWSON¹⁴ ; J. PETERSEN¹ ; J. SCHLERETH¹⁴ ; J. STRONG⁶ ; J. VERMEULEN¹¹ ; J.A. BOGAERTS¹ ; K. KORCYL³ ; K. NAKAYOSHI¹⁵ ; L. MAPELLI¹ ; L. TREMBLET¹ ; M. ABOLINS¹⁶ ; M. BERETTA¹⁷ ; M. CIOBOTARU¹ ; M. GRUWE¹ ; M. JOOS¹ ; M. LE VINE¹⁸ ; M. LOSADA MAIA² ; M. MULLER⁴ ; M. SHIMOJIMA¹⁹ ; M. YU⁴ ; M. ZUREK³ ; M.L. FERRER²⁰ ; P. DE MATOS LOPES PINTO¹ ; P. GOLONKA¹ ; P. WERNER¹ ; R. BLAIR¹⁴ ; R. CRANFIELD¹⁰ ; R. DOBINSON¹ ; R. FERRARI²¹ ; R. HAUSER¹⁶ ; R. Hughes-Jones²² ; R. MCLAREN¹ ; R. MOMMSEN⁵ ; R. SPIWOKS¹ ; S. GADOMSKI⁷ ; S. GAMEIRO¹ ; S. HAAS¹ ; S. STANCU¹ ; T. MAENO¹ ; V. PEREZ REALE⁷ ; W. VANDELLI²¹ ; Y. ERMOLINE¹⁶ ; Y. HASEGAWA²³ ; Y. NAGASAKA²⁴ ; Y. YASU¹⁵

¹ CERN, Geneva, Switzerland

² Universidade Federal do Rio de Janeiro, COPPE/EE, Brazil

³ Polish Academy of Sciences, Cracow, Poland

⁴ Universitaet Mannheim, Germany

⁵ University of California, Irvine, USA

⁶ University of London, Egham, UK

⁷ University of Bern, Switzerland

⁸ RAL, Didcot, UK

⁹ INFN Roma, Italy

¹⁰ University College London, UK

¹¹ NIKHEF, The Netherlands

¹² UNIVERSITY OF CALIFORNIA AT IRVINE AND CERN

¹³ University of Wisconsin, USA

¹⁴ Argonne National Laboratory, USA

¹⁵ KEK, Japan

¹⁶ Michigan State University, USA

¹⁷ INFN Frascati, Italy

¹⁸ BNL, USA

¹⁹ Nagasaki Institute of Applied Science, Japan

²⁰ INFN Frascati, Italy

²¹ INFN Pavia, Italy

²² THE UNIVERSITY OF MANCHESTER

²³ Shinshu University, Japan

²⁴ Hiroshima Institute of Technology, Japan

Auteur(s) contact: gokhan.unel@cern.ch

The ATLAS Trigger and DAQ system is designed to use the Region of Interest (RoI) mechanism to reduce the initial Level 1 trigger rate of 100 kHz down to about 3.3 kHz Event Building rate.

The DataFlow component of the ATLAS TDAQ system is responsible for the reading of the detector specific electronics via 1600 point to point readout links, the collection and provision of RoI to the Level 2 trigger, the building of events accepted by the Level 2 trigger and their subsequent input to the Event Filter system where they are subject to further selection criteria.

To validate the design and implementation of the DAQ DataFlow system, a prototype setup representing 20% of the final system, has been put together at CERN. This baseline prototype contains 68 PCs running Linux, and exchanging data via a 64-port and a 31-port Gigabit Ethernet switches for Event Building and RoI Collection. The system performance is measured by playing back simulated data through the system and running prototype algorithms in the Level 2 trigger. In parallel a full discrete event model of the system has been developed and tuned to the testbed results as an aid to studying the system performance at and beyond the size of the prototype setup.

Measurements will be presented on the performance of the prototype setup, showing that the components of the current integrated system implementation can already sustain their nominal ATLAS requirements using existing hardware and Gigabit network technology: 20 kHz RoI Collection rate per readout link, 3 kHz Event Building rate and 70 Mbyte/s throughput per event building node. The use of these results to calibrate the model will also be presented along with the model predictions for the performance of the final DAQ DataFlow system.

Distributed Computing Systems and Experiences / 499

Performance of the NorduGrid ARC and the Dulcinea Executor in ATLAS Data Challenge 2

A. Filipic¹; A. Konstantinov²; A. Read³; A. Sandgren⁴; A. Teras⁵; A. Waananen⁶; B. Konya⁷; C. Haerberli⁸; D. Kalici⁹; F. Ould-Saada¹⁰; G. Duckeck¹¹; H. Heller¹²; H. Jensen¹³; H. Riiser³; J. R. Hansen⁶; J. Kennedy¹¹; J. Kleist¹³; J. Lindemann¹⁴; J.F. Mycklebust¹⁵; J.L. Nielsen⁶; K. Pajchel³; L. Nixon¹⁶; L. malinowsky¹⁷; M. Ellert¹⁸; M. Hanshaugen¹⁹; M. Tadel¹; M. Wadenstein⁴; N. Smeds¹⁷; O. Smirnova²⁰; P. Eerola⁷; R. Sturrock²¹; S. Sunde¹⁹; T. Ekelof²²

¹ JSI Ljubljana

² Vilnius University, Lithuania

³ University of Oslo

⁴ Umea University

⁵ Helsinki

⁶ Niels Bohr Institute, Copenhagen, Denmark

⁷ Lund University, Sweden

⁸ University of Bern

⁹ Odense

¹⁰ University of Oslo, Norway

¹¹ LMU Munich

¹² LR, Munich

¹³ Alborg University

¹⁴ Lund

¹⁵ Bergen

¹⁶ Linkoping University

¹⁷ KTH Stockholm

¹⁸ CERN

¹⁹ USIT Oslo

²⁰ LUND UNIVERSITY

²¹ University of Melbourne

²² Uppsala University, Sweden

Auteur(s) contact: oxana.smirnova@hep.lu.se

This talk describes the various stages of ATLAS Data Challenge 2 (DC2) in what concerns usage of resources deployed via NorduGrid's Advanced Resource Connector (ARC). It also describes the integration of these resources with the ATLAS production system using the Dulcinea executor.

ATLAS Data Challenge 2 (DC2), run in 2004, was designed to be a step forward in the distributed data processing. In particular, much coordination of task assignment to resources was planned to be delegated to Grid in its different flavours. An automatic production management system was designed, to direct the tasks to Grids and conventional resources.

The Dulcinea executor is a part of this system that provides interface to the information system and resource brokering capabilities of the ARC middleware. The executor translates the job definitions received from the supervisor to the extended resource specification language (XRSL) used by the ARC middleware. It also takes advantage of the ARC middleware's built-in support for the Globus Replica Location Server (RLS) for file registration and lookup.

NorduGrid's ARC has been deployed on many ATLAS-dedicated resources across the world in order to enable effective participation in ATLAS DC2. This was the first attempt to harness large amounts of strongly heterogeneous resources in various countries for a single collaborative exercise using Grid tools. This talk addresses various issues that arose during different stages of DC2 in this environment: preparation, such as ATLAS software installation; deployment of the middleware; and processing. The results and lessons are summarized as well.

Poster Session 3 - Tableau: 22 / 427

Persistence for Analysis Objects

J. Hrivnac¹

¹ LAL

Auteur(s) contact: julius.hrivnac@cern.ch

There are two kinds of analysis objects with respect to their persistent requirements:
Objects, which need direct access to the persistency service only for their IO operations (read/write/update/...): histograms, clouds, profiles, ...
All Persistency requirements for those objects can be implemented by standard Transient-Persistent Separation techniques like JDO, Serialization, etc.
 Objects, which need direct access to the persistency service for

some of their standard operations: NTuples, Tags,... It is not feasible to completely separate Transient and Persistent form of those objects. Their Persistency should be tightly interfaced with their transient form. One possibility is to directly implement a persistent extension of those objects for each persistency mechanism. The SQLTuple has been developed to deliver efficient SQL persistency for AIDA standard NTuple objects. The implementation is based on FreeHEP AIDA implementation and is completely inter-operable with other FreeHEP components as well as with other AIDA implementations. SQLTuple dependency on SQL database implementation is handled at run-time by textual configuration. In principle all mainstream SQL databases are supported. The default mapping layer can be customized so that, for example, LCG Pool Tag databases can be transparently supported. This customization is used to implement higher level management utilities for Pool Tag databases - package ColMan. ColMan utilities are accessible also from the C++ environment and via standard Web Service. The representation will cover both SQLTuple and ColMan packages and their inter-operability with other tools. Performance assessment of various available technologies will be covered as well.

Plenary / 510

Physics Validation of the LHC Software

F. GIANOTTI^{None}

The LHC Software will be confronted to unprecedented challenges as soon as the LHC will turn on. We summarize the main Software requirements coming from the LHC detectors, triggers and physics, and we discuss several examples of Software components developed by the experiments and the LCG project (simulation, reconstruction, etc.), their validation, and their adequacy for LHC physics.

Event Processing / 493

Physics validation of the simulation packages in a LHC-wide effort

A. Ribon¹

¹ CERN

Auteur(s) contact: alberto.ribon@cern.ch

In the framework of the LCG Simulation Physics Validation Project, we present comparison studies between the GEANT4 and FLUKA shower packages and LHC sub-detector test-beam data. Emphasis is given to the response of LHC calorimeters to electrons, photons, muons and pions. Results of “simple-benchmark” studies, where the above simulation packages are compared to data from nuclear facilities, are also shown.

Poster Session 3 - Tableau: 14 / 378

Pixel Reconstruction in the CMS High-Level Trigger

D. Kotlinski¹ ; M. Konecki² ; S. Cucciarelli³ ; T. Todorov⁴

¹ *Paul Scherrer Institut (PSI)*

² *University of Basel*

³ *CERN*

⁴ *Universite Louis Pasteur*

Auteur(s) contact: susanna.cucciarelli@cern.ch

The Pixel Detector is the innermost one in the tracking system of the Compact Muon Solenoid (CMS) experiment. It provides the most precise measurements not only supporting the full track reconstruction but also allowing the standalone reconstruction useful especially for the online event selection at High-Level Trigger (HLT). The performance of the Pixel Detector is given. The HLT algorithms using Pixel Detector are presented, including pixel track reconstruction, primary vertex finding, tau identification, isolation and track seeding.

Poster Session 1 - Tableau: 4 / 251

Portable Gathering System for Monitoring and Online Calibration at Atlas

ATLAS TDAQ HLT group^{None} ; P. Conde MUINO¹

¹ *CERN*

Auteur(s) contact: patricia.conde.muino@cern.ch

During the runtime of any experiment, a central monitoring system that detects problems as soon as they appear has an essential role. In a large experiment, like Atlas, the online data acquisition system is distributed across the nodes of large farms, each of them running several processes that analyse a fraction of the events. In this architecture, it is necessary to have a central process that collects all the monitoring data from the different nodes, produces full statistics histograms and analyses them. In this paper we present the design of such a system, called the “gatherer”. It allows to collect any monitoring object, such as histograms, from the farm nodes, from any process in the DAQ, trigger and reconstruction chain. It also adds up the statistics, if required, and processes user defined algorithms in order to analyse the monitoring data. The results are sent to a centralized display, that shows the information online, and to the archiving system, triggering alarms in case of problems.

The innovation of our approach is that conceptually it abstracts the several communication protocols underneath, being able to talk with different processes using different protocols at the same time and, therefore, providing maximum flexibility. The software is easily adaptable to any trigger-DAQ system.

The first prototype of the gathering system has been implemented for Atlas and will be running during this year’s combined test beam. An evaluation of this first prototype will also be presented.

Poster Session 3 - Tableau: 28 / 183

Porting CLEO software to Linux

C. Jones¹ ; V. Kuznetsov¹

¹ CORNELL UNIVERSITY

Auteur(s) contact: valentin.kuznetsov@cern.ch

Linux operating system has become the platform of choice in the HEP community. However, the migration process from another operating system to Linux can be a tremendous effort for developers and system administrators. The ultimate goal of such a transition is to maximize agreement between the final results of identical calculations on the different platforms. Apart from the fine tuning of the existing software the following issues need to be resolved: choice of Linux distribution, development tools (compiler, debugger, profilers etc.), compatibility with 3d party software, and deployment strategy. It would be ideal to develop, run and test software using office desktops, local farm systems, or personal laptop regardless of the Linux distribution choosen. To accomplish this task you need to have flexible package management system which is capable to install/upgrade/verify/uninstall necessary software components without particular knowledge of remote system configuration and user privileges. We discuss how Linux became the third official computing platform of the CLEO collaboration, outlining the details of the transition from OSF and Solaris operating systems to Linux, software model and deployment strategy employed.

Poster Session 3 - Tableau: 48 / 257

Porting LCG Applications

I. Reguero¹ ; J A. Lopez-Perez¹

¹ CERN, IT DEPARTMENT

Auteur(s) contact: ignacio.reguero@cern.ch

Our goal is two fold. On one hand we wanted to address the interest of CMS users to have LCG Physics analysis environment on Solaris. On the other hand we wanted to assess the difficulty of porting code written in Linux without particular attention to portability to other Unix implementations. Our initial assumption was that the difficulty would be manageable even for a very small team. This is because the implicit respect by Linux of most Unix interfaces and standards such as the IEEE (PASC) 1003.1 1003.2 specifications.

We started with the LCG External software (<http://spi.web.cern.ch/spi/extsoft/platform.html>) in order to use it to build the LCG applications such as POOL and SEAL (<http://lcgapp.cern.ch/project/>).

We will discuss the main problems found with the system interfaces as well as the advantages and disadvantages of using the GNU compilers and development environment versus the vendor provided ones.

Distributed Computing Services / 420

Practical approaches to Grid workload and resource management in the EGEE project

A. GUARISE¹ ; A. Gianelle² ; A. KRENEK³ ; A. MARASCHINI⁴ ; A. TERRACINA⁴ ; D. KOURIL³ ; D. REBATTO⁵ ; E. RONCHIERI⁶ ; F. GIACOMINI⁶ ; F. PACINI⁴ ; F. PRELZ⁵ ; G. AVELLINO⁴ ; G. PATANIA¹ ; J. POSPISIL³ ; J. SITERA³ ; J. SKRABAL³ ; L. MATYSKA³ ; L. ZANGRANDO⁷ ; M. MARCHI⁵ ; M. MEZZADRI⁵ ; M. MORDACCHINI⁷ ; M. MULAC³ ; M. PAPPALARDO⁸ ; M. RUDA³ ; M. Sgaravatto⁷ ; M. VOCU³ ; P. ANDREETTO⁷ ; S. ANDREOZZI⁶ ; S. BECO⁴ ; S. BORGIA⁷ ; S. MONFORTE⁸ ; V. CIASCHINI⁶ ; Z. SALVET³

¹ INFN Torino

² INFN Padova, Italy

³ CESNET

⁴ DATAMAT

⁵ INFN Milano

⁶ INFN Cnaf

⁷ INFN Padova

⁸ INFN Catania

Auteur(s) contact: massimo.sgaravatto@pd.infn.it

Resource management and scheduling of distributed, data-driven applications in a Grid environment are challenging problems. Although significant results were achieved in the past few years, the development and the proper deployment of generic, reliable, standard components present issues that still need to be completely solved. Interested domains include workload management, resource discovery, resource matchmaking and brokering, accounting, authorization policies, resource access, reliability and dependability. The evolution towards a service-oriented architecture, supported by emerging standards, is another activity that will demand attention.

All these issues are being tackled within the EU-funded EGEE project (Enabling Grids for E-science in Europe), whose primary goals are the provision of robust middleware components and the creation of a reliable and dependable Grid infrastructure to support e-Science applications.

In this paper we present the plans and the preliminary activities aiming at providing adequate workload and resource management components, suitable to be deployed in a production-quality Grid.

Event Processing / 351

Precision electromagnetic physics in Geant4: the atomic relaxation models

A. Mantero¹ ; M.G. Pia² ; S. Guatelli³ ; S. Saliceti¹

¹ INFN Genova

² INFN GENOVA

³ INFN Genova, Italy

Auteur(s) contact: maria.grazia.pia@cern.ch

Various experimental configurations - such as, for instance, some gaseous detectors, require a high precision simulation of electromagnetic physics processes, accounting not only for the primary interactions of particles with matter, but also capable of

describing the secondary effects deriving from the de-excitation of atoms, where primary collisions may have created vacancies. The Geant4 Simulation Toolkit encompasses a set of models to handle the atomic relaxation induced by the photoelectric effect, Compton scattering and ionization, with the production of X-ray fluorescence and of Auger electrons. We describe the physics models implemented in Geant4 to handle the atomic relaxation, the object-oriented design of the software and the validation of the models with respect to test beam data. In particular, we present a novel development of an original model for particle induced X-ray emission, to be released for the first time in the summer of 2004. We illustrate applications of Geant4 atomic relaxation models for physics reach studies in a real-life experimental context

Poster Session 3 - Tableau: 44 / 342

Precision validation of Geant4 electromagnetic physics

Auteur(s): M.G. Pia¹

Co-auteur(s) A. Mantero ²; B. Mascialino ²; L. Pandola ³; M. Maire ⁴; P. Cirrone ⁵; S. Donadio ²; S. Guatelli ²; S. Parlati ³; V. Ivanchenko ⁶

¹ INFN GENOVA

² INFN Genova

³ INFN LNGS

⁴ LAPP

⁵ INFN LNS

⁶ Budker Institute

Auteur(s) contact: maria.grazia.pia@cern.ch

The Geant4 Toolkit provides an ample set of alternative and complementary physics models to handle the electromagnetic interactions of leptons, photons, charged hadrons and ions. Because of the critical role often played by simulation in the experimental design and physics analysis, an accurate validation of the physics models implemented in Geant4 is essential, down to the quantitative understanding of the accuracy of their microscopic features. Results from a series of detailed tests with respect to well established reference data sources and experiments are presented, focusing in particular on the precision validation of the microscopic components of Geant4 physics, such as cross sections and angular distributions, provided in the various alternative physics models of Geant4 electromagnetic packages. The validation of Geant4 physics is performed by means of quantitative evaluations of the comparison of Geant4 models to reference data are presented, making use of statistical analysis algorithms to estimate the compatibility of simulated and experimental distributions. Such precision tests are especially relevant for critical applications of simulation models, such as tracking detectors, neutrino and other astroparticle experiments, medical physics

Poster Session 2 - Tableau: 2 / 273

Predicting Resource Requirements of a Job Submission

A. Ali¹ ; A. Anjum¹ ; A. Mehmood¹ ; H. Newman² ; I. Willers² ; J. Bunn² ; R. Meclatchey³

¹ NIIT

² Caltech

³ UWE

Auteur(s) contact: ashiq.anjum@niit.edu.pk

Grid computing provides key infrastructure for distributed problem solving in dynamic virtual organizations. However, Grids are still the domain of a few highly trained programmers with expertise in networking, high-performance computing, and operating systems.

One of the big issues in the full-scale usage of a grid is the matching of the resource requirements of a job submission to available resources. In order for resource brokers/job schedulers to ensure efficient use of grid resources, an initial estimate of the likely resource usage of a submission must be made. In the context of the Grid Enabled Analysis Environment (GAE), physicists want the ability to discover, acquire, and reliably manage computational resources dynamically, in the course of their everyday activities. They do not want to be bothered with the location of these resources, the mechanisms that are required to use them, keeping track of the status of computational tasks operating on these resources, or with reacting to failure. They do care about how long their tasks are likely to run and how much these tasks will cost.

So the grid scheduler must have the capability to estimate before job submission, how much time and resources the job will consume on execution site. Our proposed module, Prediction engine will be part of scheduler and it will provide estimates of resource use along with the duration of use. This will enable scheduler to choose the optimum site for job execution.

This paper presents the survey of existing grid schedulers and then based on this survey states the need for resource usage estimation. Also the architecture and design of “grid prediction engine” that predicts the resource requirements of a job submission is discussed.

Poster Session 2 - Tableau: 29 / 245

Production Experience of the Storage Resource Broker in the BaBar Experiment

A. Hasan¹ ; J-Y. Nief² ; W. Kroeger¹

¹ SLAC

² ccin2p3

Auteur(s) contact: hasan@slac.stanford.edu

We describe the production experience gained from implementing and using exclusively the San Diego Super Computer Center developed Storage Resource Broker (SRB) to distribute the BaBar experiment’s production event data stored in ROOT files from the experiment center at SLAC, California, USA to a Tier A computing center at ccin3, Lyon France. In addition we outline how the system can be readily expanded to include more sites.

Distributed Computing Systems and Experiences / 231

Production Management Software for the CMS Data Challenge

A. Anzar¹ ; A. Fanfani² ; C. Grandi³ ; G. Graham¹ ; G. Guglielmo¹ ; J. Andreeva⁴ ; N. Darmanov⁵ ; N. Ratnikova¹ ; P. GARCIA-ABIA⁶ ; T. Wildish⁷ ; V. Lefebure⁵

¹ *Fermilab*

² *INFN, Bologna*

³ *INFN-Bologna*

⁴ *UC Riverside*

⁵ *CERN*

⁶ *CIEMAT*

⁷ *Princeton University*

Auteur(s) contact: julia.andreeva@cern.ch

One of the goals of CMS Data Challenge in March-April 2004 (DC04) was to run reconstruction for sustained period at 25 Hz input rate with distribution of the produced data to CMS T1 centers for further analysis.

The reconstruction was run at the T0 using CMS production software, of which the main components are RefDB (CMS Monte Carlo 'Reference Database' with Web interface) and McRunjob (a framework for creation and submission of large numbers of Monte Carlo jobs).

This paper presents an overview of CMS production cycle, describing production tools, covering data processing, bookkeeping and publishing issues, in the context of their use during the T0 reconstruction part of DC04.

Poster Session 2 - Tableau: 30 / 332

Production data export and archiving system for new data format of the BaBar experiment.

A. Hasan¹ ; A. Trunov¹ ; T. Azemoun¹ ; W. Kroeger¹

¹ *SLAC*

Auteur(s) contact: artem@slac.stanford.edu

For The BaBar Computing Group

BaBar has recently moved away from using Objectivity/DB for its event store towards a ROOT-based event store. Data in the new format is produced at about 20 institutions worldwide as well as at SLAC. Among new challenges are the organization of data export from remote institutions, archival at SLAC and making the data visible to users for analysis and import to their own institutions.

The new system is designed to be scalable, easily configurable on the client and server side and adaptive to server load. It's integrated to work with SLAC's mass storage system (HPSS) and with the xrootd service. Design, implementation and experience with new system, as well as future development is discussed in this article.

Distributed Computing Services / 334

Production mode Data-Replication framework in STAR using the HRM Grid

A. Shoshani¹ ; A. Sim¹ ; D. Olson¹ ; E. Hjort² ; J. Lauret³

¹ *Lawrence Berkeley Laboratory*

² *LAWRENCE BERKELEY LABORATORY*

³ *Brookhaven National Laboratory*

Auteur(s) contact: elhjort@lbl.gov

The STAR experiment utilizes two major computing facilities for its data processing needs - the RCF at Brookhaven and the PDSF at LBNL/NERSC. The sharing of data between these facilities utilizes data grid services for file replication, and the deployment of these services was accomplished in conjunction with the Particle Physics Data Grid (PPDG). For STAR's 2004 run it will be necessary to replicate ~100 TB. The file replication is based on Hierarchical Resource Managers (HRMs) along with Globus tools for security (GSI) and data transport (GridFTP). HRMs are grid middleware developed by the Scientific Data Management group at LBNL, and STAR file replication consists of an HRM interfaced to HPSS at each site with GridFTP transfers between the HRMs. Each site also has its own installation of the STAR file and metadata catalog, which is implemented in MySQL. Queries to the catalogs are used to generate file transfer requests. Single requests typically consist of many thousands of files with a volume of hundreds of GBs. The HRMs implement a plugin to a Replica Registration Service (or RRS) which is utilized for automatic registration of new files as they are successfully transferred across sites. This allows STAR users immediate use of the distributed data. Data transfer statistics and system architecture will be presented.

Poster Session 2 - Tableau: 10 / 129

Production of simulated events for the BaBar experiment by using LCG

C. Bozzi¹ ; D. Andreotti² ; E. Antonioli² ; E. Luppi² ; M. Melani³ ; P. Veronesi²

¹ *INFN SEZIONE DI FERRARA*

² *INFN Sezione di Ferrara*

³ *SLAC*

Auteur(s) contact: bozzi@fe.infn.it

The BaBar experiment has been taking data since 1999. In 2001 the computing group started to evaluate the possibility to evolve toward a distributed computing model in a Grid environment. In 2003, a new computing model, described in other talks, was implemented, and ROOT I/O is now being used as the Event Store. We implemented a system, based on the LHC Computing Grid (LCG) tools, to submit full-scale MonteCarlo simulation jobs in this new BaBar computing model framework. More specifically, the resources of the LCG implementation in Italy, grid.it, are used as computing elements (CE) and Worker Nodes (WN). A Resource Broker (RB) specific for the BaBar computing needs was installed. Other BaBar requirements, such as the installation and usage of an object-oriented (Objectivity) Database to read detector conditions and calibration constants, were accommodated by using non-gridified hardware in a subset of grid.it sites. The BaBar simulation software was packed and installation on Grid elements was centrally managed with LCG tools. Sites were geographically mapped to Objectivity databases, and conditions were read by the WN either locally or remotely. An LCG User Interface (UI) has been used to submit simulation tests by using standard JDL commands. The ROOT I/O output files were retrieved from the WN and stored in the closest Storage Element (SE). Standard BaBar simulation production tools were then installed on the UI and configured such that the resulting simulated events can be merged and shipped to SLAC, like in the standard BaBar simulation production setup. Final validation of the system is being completed. This gridified approach results in

the production of simulated events on geographically distributed resources with a large throughput and minimal, centralized system maintenance.

Poster Session 3 - Tableau: 29 / 34

PyBus – A Python Software Bus

W. Lavrijsen¹

¹ *LBNL*

Auteur(s) contact: wim.lavrijsen@cern.ch

A software bus, just like its hardware equivalent, allows for the discovery, installation, configuration, loading, unloading, and run-time replacement of software components, as well as channeling of inter-component communication. Python, a popular open-source programming language, encourages a modular design on software written in it, but it offers little or no component functionality. However, the language and its interpreter provide sufficient hooks to implement a thin, integral layer of component support. This functionality can be presented to the developer in the form of a module, making it very easy to use. This paper describes a Python module, PyBus, with which the concept of a ‘software bus’ can be realised in Python. It demonstrates, within the context of the Atlas software framework Athena, how PyBus can be used for the installation and (run-time) configuration of software, not necessarily Python modules, from a Python application in a way that is transparent to the end-user.

Event Processing / 379

Python-based physics analysis environment for LHCb

A. TSAREGORODTSEV¹ ; E. DE OLIVEIRA² ; G B. BARRAND³ ; I. BELYAEV⁴ ; M. FRANK⁵ ; P. MATO⁵

¹ *Uni Marseil II/Marseil*

² *CBPF/Rio de Janeiro*

³ *LAL/Orsay*

⁴ *ITEP/Moscow*

⁵ *CERN*

Auteur(s) contact: pere.mato@cern.ch, ivan.belyaev@itep.ru

Bender, the Python based physics analysis application for LHCb combines the best features of underlying Gaudi C++ software architecture with the flexibility of Python scripting language and provides end-users with friendly physics analysis oriented environment. It is based in one hand, on the generic Python bindings for the Gaudi framework, called GaudiPython, and in the other hand on an efficient C++ physics analysis toolkit called LoKi. Bender and LoKi use the tools from the physics analysis framework, called DaVinci. Bender achieves a clear separation between the technical details and the physical contents of end-user physicist’s code. The usage of Python, AIDA abstract interfaces and standard LCG reflection techniques allows an easy integration of Bender’s analysis environment with third party products like the interactive event display and visualization tools like Panoramix/LaJoconde, ROOT and HippoDraw. We’ll present the overall design and capabilities of the system, its status and prospects.

Core Software / 361**Quality Assurance and Testing in LCG**A. Aimar¹ ; J. Moscicki¹ ; L. MANCERA¹ ; M. GALLAS¹ ; M. Lamanna¹¹ CERN

Software Quality Assurance is an integral part of the software development process of the LCG Project and includes several activities such as automatic testing, test coverage reports, static software metrics reports, bug tracker, usage statistics and compliance to build, code and release policies.

As a part of QA activity all levels of the sw-testing should be run as a part of automatic process, the SPI project delivers a general test-framework solution based on open source software together with test document templates and software testing policies. The test-framework solution is built on QMtest, Oval and the X-Unit family (CppUnit, PyUnit, JUnit). The specific languages testing features are covered at the unit-testing level with the X-Unit family, the validation testing activity can be done through Oval. And Qmtest offers a way to integrate all the tests and write custom python tests, having a nice web interface for running and browse the test results.

Test coverage reports allow to understand to which extent software products are tested and they are based on the approach used by Linux Testing Project. Code size and development effort of the software is estimated using sloccount utility based on standard development models. Statistics are automatically extracted from the savannah bug tracker which enables to analyze the evolution of the quality, amount of feedback from the users etc. Finally the compliance with the standard LCG policies is verified. It includes the build and CVS repository structure and the standard release procedure.

Poster Session 3 - Tableau: 39 / 98**RDBC: ROOT DataBase Connectivity****Auteur(s):** V. Onuchin¹**Co-auteur(s)** Fons Rademakers² ; P. Canal³ ; Rene Brun²¹ CERN, IHEP² CERN³ FNAL**Auteur(s) contact:** valeri.onouchine@cern.ch

The RDBC (ROOT DataBase Connectivity) library is a C++ implementation of the The Java Database Connectivity Application Programming Interface. It provides a DBMS-independent interface to relational databases from ROOT as well as a generic SQL database access framework.

RDBC also extends the ROOT TSQL abstract interface.

Currently it is used in two large experiments:

- in Minos as interface to MySQL and Oracle databases
- in Phenix as interface to PostgreSQL database.

In this paper we will describe the main features and applicability of this library.

Poster Session 3 - Tableau: 4 / 298

ROOT : detector visualization

A. Gheata¹ ; M. Gheata² ; O. Couet¹ ; R. Brun¹

¹ CERN

² Institute of Space Science - Bucharest

Auteur(s) contact: olivier.couet@cern.ch

The ROOT geometry package is a tool designed for building, browsing, tracking and visualizing a detector geometry. The code is independent from other external MC for simulation, therefore it does not contain any constraints related to physics. However, the package defines a number of hooks for tracking, such as media, materials, magnetic field or track state flags, in order to allow interfacing to tracking MC's. The final goal is to be able to use the same geometry for several purposes, such as tracking, reconstruction or visualization, taking advantage of the ROOT features related to bookkeeping, I/O, histograming, browsing and GUI's.

In this poster, we will show the various graphics tools to render complex geometries, from ray tracing tools that have the advantage to test the real geometry like when tracking particles, to sophisticated 3-D dynamic graphics with the OpenGL, X3D, Coin3D or OpenInventor viewers. An abstract interface has been defined and it is common to all the viewers.

Poster Session 3 - Tableau: 3 / 170

ROOT Graphical User Interface

B. Bellenot¹ ; F. Rademakers² ; I. Antcheva² ; O. Couet² ; R. Brun² ; V. Onuchin²

¹ Alcan

² CERN

Auteur(s) contact: ilka.antcheva@cern.ch

The GUI is a very important component of the ROOT framework. Its main purpose is to improve the usability and end-user perception. In this paper, we present two main projects in this direction: the ROOT graphics editor and the ROOT GUI builder.

The ROOT graphics editor is a recent addition to the framework. It provides a state of the art and an intuitive way to create or edit objects in the canvas.

The ROOT GUI builder greatly facilitates the design, the development and the maintenance of any interactive application based on the ROOT framework. GUI objects can be selected, dragged/dropped in the widgets. An automatic code generator can be activated to save the code corresponding to any complex layout. This code can be executed via the CINT interpreter or directly compiled with the user application.

Past surveys indicate that the development of a GUI is a significant undertaking and that the GUI's source code is a substantial portion of

the program's overall source base. The new GUI builder in ROOT will enable the rapid construction of simple and complex GUIs.

Poster Session 1 - Tableau: 6 / 84

Raw Ethernet based hybrid control system for the automatic control of suspended masses in gravitational waves interferometric detectors

A. Eleuteri¹ ; F. Acernese² ; F. Barone³ ; G. RUSSO⁴ ; I. Ricciardi⁴ ; K. Qipiani⁵ ; L. Milano⁴ ; R. De Rosa⁴ ; S. Pardi⁶

¹ *DIPARTIMENTO DI SCIENZE FISICHE - UNIVERSITÀ DI NAPOLI FEDERICO II*

² *Dipartimento di Scienze Fisiche - Università di Salerno*

³ *Dipartimento di Scienze Farmaceutiche - Università di Salerno*

⁴ *Dipartimento di Scienze Fisiche - Università di Napoli Federico II*

⁵ *INFN - sez. Napoli*

⁶ *Dipartimento di Matematica ed Applicazioni "R.Caccioppoli"*

Auteur(s) contact: antonio.eleuteri@na.infn.it

In this paper we examine the performance of the raw Ethernet protocol in deterministic, low-cost, real-time communication. Very few applications have been reported until now, and they focus on the use of the TCP and UDP protocols, which however add a sensible overhead to the communication and reduce the useful bandwidth. We show how low-level Ethernet access can be used for peer-to-peer, short distance communication, and how it allows the writing of applications requiring large bandwidth. We show some examples running on the Lynx real-time OS and on Linux, both in mixed and homogeneous environments. As an example of application of this technique, we describe the architecture of an hybrid Ethernet based real-time control system prototype we implemented in Napoli, discussing its characteristics and performances. Finally we discuss its application to the real-time control of a suspended mass of the mode cleaner of the 3m prototype optical interferometer for gravitational wave detection operational in Napoli.

Wide Area Networking / 295

Realization of a stable network flow with high performance communication in high bandwidth-delay product network

Auteur(s): Y. Kodama¹

Co-auteur(s) O. Tatebe² ; S. Sekiguchi² ; T. Kudoh²

¹ *NATIONAL INSTITUTE OF ADVANCED INDUSTRIAL SCIENCE AND TECHNOLOGY (AIST)*

² *AIST*

Auteur(s) contact: y-kodama@aist.go.jp

It is important that the total bandwidth of the multiple streams should not exceed the network bandwidth in order to achieve a stable network

flow with high performance in high bandwidth-delay product networks. Software control of bandwidth for each stream sometimes exceed the specified bandwidth. We proposed the hardware control technique for total bandwidth of multiple streams with high accuracy.

GNET-1 is the hardware gigabit network testbed that we developed. It provides functions such as wide area network emulation, network instrumentation, and traffic generation at gigabit Ethernet wire speeds. GNET-1 is a powerful tool for developing network-aware grid software. It can control the total bandwidth of the multiple streams with high accuracy by adjusting the interframe gap (IFG).

To see the effect of the highly accurate bandwidth control by GNET-1, the file exchange of large-scale data was done on a Trans-pacific Grid Datafarm testbed between Japan-U.S.. We used three trans-pacific networks, APAN/TransPAC Los Angeles line and its Chicago line and SuperSINET New York line. Its total bandwidth that can be used was 3.9 Gbps. In this feasible study, GNET-1 controlled five gigabit Ethernet ports, and achieved the total bandwidth of 3.78 Gbps in stable for about one hour. The bandwidth was 97 % of the peak bandwidth of used networks.

Core Software / 85

Recent Developments in the ROOT I/O

F. Rademakers¹ ; P. Canal² ; R. Brun¹

¹ CERN

² FERMILAB

Auteur(s) contact: philippe.canal@cern.ch

Since version 3.05/02, the ROOT I/O System has gone through significant enhancements.

In particular, the STL container I/O has been upgraded to support splitting, reading without existing libraries and using directly from TTreeFormula (TTree queries).

This upgrade to the I/O system is such that it can be easily extended (even by the users) to support the splitting and querying of almost any collections. The ROOT TTree queries engine has also been enhanced in many ways including an increase performance, better support for array printing and histogramming, addition of the ability to call any external C or C++ functions, etc.

We improved the I/O support for classes not inheriting from TObject, including support for automatic schema evolution without using an explicit class version. ROOT now support generating files larger than 2Gb. We also added plugins for several of the mass storage servers (Castor, DCache, Chirp, etc.).

We will describe in details these new features and their implementation.

Poster Session 3 - Tableau: 5 / 174

Recent evolutions of CMT. Multi-project and activity management.

C. ARNAULT¹

¹ CNRS

Auteur(s) contact: arnault@lal.in2p3.fr

Since its introduction in 1999, CMT is now used as a production tool in many large software projects for physics research (ATLAS, LHCb, Virgo, Auger, Planck). Although its basic concepts remain unchanged since the beginning, proving their viability, it is still improving and increasing its coverage of the configuration management mechanisms. Two important evolutions have recently been introduced, one for explicitly supporting multi-project environments, and the other to specify and manage configuration activities.

The existing concept of package area is now extended to cover the support of sub-projects structuring, with the possibility of assigning configuration management properties (typically strategies) to each sub project, allowing for instance to have installation area mechanisms only applicable for some of them.

It is also possible to specify parameterized activities that will be run on demand either through make or through an explicit activation command, which ensures that the runtime environment is properly setup.

Core Software / 494

Reflection-Based Python-C++ Bindings

J. Generowicz¹ ; M. Marino² ; P. Mato¹ ; W. LAVRIJSEN²

¹ CERN

² LBNL

Auteur(s) contact: wim.lavrijsen@cern.ch, pere.mato@cern.ch

Python is a flexible, powerful, high-level language with excellent interactive and introspective capabilities and a very clean syntax. As such it can be a very effective tool for driving physics analysis.

Python is designed to be extensible in low-level C-like languages, and its use as a scientific steering language has become quite widespread. To this end, existing and custom-written C or C++ libraries are bound to the Python environment as so-called extension modules. A number of tools for easing the process of creating such bindings exist, such as SWIG or Boost.Python. Yet, the the process still requires a considerable amount of effort and expertise.

The C++ language has little built-in introspective capabilities, but tools such as LCGDict and CINT add this by providing so-called dictionaries: libraries that contain information about the names, entry points, argument types, etc. of other libraries.

The reflection information from these dictionaries can be used for the creation of bindings and so the process can be fully automated, as dictionaries are already provided for many end-user libraries for other purposes, such as object persistency.

PyLCGDict is a Python extension module that uses LCG dictionaries, as PyROOT uses CINT reflection information, to allow Python users to access C++ libraries with essentially no preparation on the users' behalf. In addition, and in a similar way, PyROOT gives ROOT users access to Python libraries.

Poster Session 1 - Tableau: 26 / 313

Remote Shifting at the CLEO Experiment

H. Schwarthoff¹ ; T. Wilksen²

¹ *Cornell University*

² *CORNELL UNIVERSITY*

Auteur(s) contact: wilksen@mail.lns.cornell.edu

The CLEO III data acquisition was from the beginning in the late 90's designed to allow remote operations and monitoring of the experiment. Since changes in the coordination and operation of the CLEO experiment two years ago enabled us to separate tasks of the shift crew into an operational and a physics task, existing remote capabilities have been revisited. In 2002/03 CLEO started to deploy its remote monitoring tasks for performing remote shifts and evaluated various communication tools e.g. video conferencing and remote desktop sharing. Remote, collaborating institutions were allowed to perform the physicist shift part from their home institutions keeping only the professional operator of the CLEO experiment on site. After a one year long testing and evaluation phase the remote shifting for physicists is now in production mode.

This talk reports on experiences made when evaluating and deploying various options and technologies used for remote control, operation and monitoring e.g. CORBA's IIOP, X11 and VNC in the CLEO experiment. Furthermore some aspects of the usage of video conferencing tools by distributed shift crews are being discussed.

Distributed Computing Services / 69

Resource Predictors in HEP Applications

J. Huth¹ ; J. Schopf² ; S. Grinstein¹

¹ *HARVARD UNIVERSITY*

² *Argonne National Lab*

Auteur(s) contact: sgrinstein@ifae.es

The ATLAS experiment uses a tiered data Grid architecture that enables possibly overlapping subsets, or replicas, of the original set to be located across the ATLAS collaboration. The full set of experiment data is located at a single Tier 0 site, and then subsets of the data are located at national Tier 1 sites, smaller subsets at smaller regional Tier 2 sites, and so on. In order to understand the data needs, both in terms of access, replication policy, and storage capacity, we need good estimations of resource needs for data manipulation. Specifically, we envision a time when a user will want to determine which is more expedient, downloading a replica from a site or recreating it from scratch.

This paper presents our technique to predict the behavior of ATLAS applications, and then to combine this information with Internet link bandwidth estimation to improve resource usage in the ATLAS Grid environment. We studied the parameters that affect the execution time performance of event generation, detector simulation, and event reconstruction. Our results show that we can achieve predictions within 10-40% of the execution time (depending on the application), better than many other pragmatic prediction techniques. We implemented a software package to provide data transfer bandwidth estimation and execution time prediction that can be used with the Chimera software to aid in managing application execution and to improve resource usage for ATLAS.

Distributed Computing Systems and Experiences / 403

Results of the LHCb experiment Data Challenge 2004

A. Pickford¹ ; A. Saroka² ; A. Tsaregorodtsev³ ; G. Kuznetsov⁴ ; G.N. Patrick⁵ ; I. Korolko⁶ ; I. Stokes-Rees² ; J. Blouw⁷ ; J. Closier⁸ ; J. Saborido-Silva⁹ ; M. Frank⁸ ; M. Gandelman¹⁰ ; M. Sanchez-Garcia⁹ ; M. Tobin¹¹ ; M. Witek¹² ; N. Brook¹³ ; P. Charpentier⁸ ; R. Bernet¹⁴ ; R. Graciani-Diaz¹⁵ ; U. Egede¹⁶ ; V. Garonne³ ; V. Romanovski¹⁷ ; V. Vagnoni¹⁸

¹ *University of Glasgow*

² *University of Oxford*

³ *Université d'Aix - Marseille II*

⁴ *Rutherford Appleton Lab, Chilton*

⁵ *Rutherford Appleton Laboratory*

⁶ *ITEP, Moscow*

⁷ *Max-Planck-Institut für Kernphysik (MPI)*

⁸ *CERN*

⁹ *Universidad de Santiago de Compostela*

¹⁰ *Universidade Federal do Rio de Janeiro (UFRJ)*

¹¹ *University of Liverpool*

¹² *High Energy Department Krakow*

¹³ *H.H. Wills Physics Laboratory*

¹⁴ *Universität Zürich*

¹⁵ *University of Barcelona*

¹⁶ *Imperial College*

¹⁷ *State Res.Center of Russian Feder. Inst.f.High Energy Phys. (IFVE)*

¹⁸ *INFN*

Auteur(s) contact: joel.closier@cern.ch

The LHCb experiment performed its latest Data Challenge (DC) in May-July 2004. The main goal was to demonstrate the ability of the LHCb grid system to carry out massive production and efficient distributed analysis of the simulation data.

The LHCb production system called DIRAC provided all the necessary services for the DC: Production and Bookkeeping Databases, File catalogs, Workload and Data Management systems, Monitoring and Accounting tools. It allowed to combine in a consistent way resources of more than 20 LHCb production sites as well as the LCG2 grid resources. 200M events constituting 90 TB of data were produced and stored in 6 Tier 1 centers. The subsequent analysis was carried out at CERN as well as in all the Tier 1 centers to where preselected datasets were distributed. The GANGA User

Interface was used to assist users in preparation of their analysis jobs and running them on the local and remote computing resources.

We will present the DC results, the experience gained utilising DIRAC and LCG2 grids as well as further developments necessary to achieve the scalability level of the real running LHCb experiment.

Poster Session 3 - Tableau: 8 / 220

Rio

G B. Barrand¹

¹ CNRS / IN2P3 / LAL

Auteur(s) contact: barrand@lal.in2p3.fr

Rio (for ROOT IO) is a rewriting of the file IO system of ROOT. We shall present our strong motivations of doing this tedious work. We shall present the main choices done in the Rio implementation (then by opposition to what we don't like in ROOT). For example, we shall say why we believe that an IO package is not a drawing package (no TClass::Draw) ; why someone should use pure abstract interfaces in such package (for example to open cleanly to various dictionaries) ; how we can have a more reliable system than ROOT (for example, by simply protect the various buffer overflows).

We shall cover the today role of Rio within OpenScientist to store histograms and tuples. We shall present the effort done around Gaudi, at the beginning of 2003, to read LHCb events with Rio (then in the "before POOL" system).

We shall present our views about the LCG proposed solution for storage, that is to say POOL over ROOT, and why the author believe that this coarse graining assembly is simply poor software engineering. We shall explain why CERN, due to its fermionic sociology, is going to miss an essential target : an appealing open source object oriented data base for HEP. We shall explain then how to do it without this lab, then passing from Rio to RioGrande...

Distributed Computing Systems and Experiences / 498

Role of Tier-0, Tier-1 and Tier-2 Regional Centres in CMS DC04

A. Anzar¹ ; A. Chierici² ; A. Fanfani³ ; A. Nowack⁴ ; A. Pacheco⁵ ; B. Martelli² ; C. Charlot⁶ ; C. Grandi⁷ ; C. Shepherd-Themistocleous⁸ ; D. Newbold⁸ ; D. Stickland⁹ ; F. Fanzago¹⁰ ; F. Martinez⁵ ; F. Rosso² ; F. Ruggieri² ; F.J. Rodriguez Calonge¹¹ ; G. Donivito¹² ; G. Lo Re¹³ ; G. Maggi¹² ; G. Merino⁵ ; G. Quast¹⁴ ; I. Fisk¹ ; I. Semeniouk⁶ ; J. Andreeva¹⁵ ; J. Hernandez¹⁶ ; J. Rehn¹⁴ ; K. Rabbertz¹⁴ ; L. Bauerdick¹ ; L. Dell'Agnello² ; L. Tuura¹⁷ ; M. Biasotto¹⁸ ; M. Corvo¹⁰ ; M. Delfino⁵ ; M. Ernst¹⁹ ; M. Mazzucato²⁰ ; M. Rodriguez⁵ ; N. Colino¹¹ ; N. De Filippis¹² ; N. Sinanis¹⁵ ; P. Capiluppi³

; P. Garcia-Abia¹¹ ; P. Mine⁶ ; P. Ricci² ; R. Harris¹ ; S. Fantinel¹⁸ ; S. Metson²¹ ; T. Barrass²¹ ; T. Wildish⁹ ; W. Jank¹⁵ ; Y. Wu¹

¹ *Fermilab*

² *INFN-CNAF*

³ *INFN Bologna*

⁴ *RWTH Aachen*

⁵ *PIC Barcelona*

⁶ *Ecole Poly*

⁷ *INFN-Bologna*

⁸ *RAL*

⁹ *Princeton University*

¹⁰ *INFN Padova*

¹¹ *CIEMAT*

¹² *INFN Bari*

¹³ *INFN & CNAF Bologna*

¹⁴ *Karlsruhe University*

¹⁵ *CERN*

¹⁶ *CIEMAT Madrid*

¹⁷ *Northeastern University*

¹⁸ *INFN Legnaro*

¹⁹ *DESY*

²⁰ *INFN-Padova*

²¹ *Bristol University*

The CMS 2004 Data Challenge (DC04) was devised to test several key aspects of the CMS Computing Model in three ways: by trying to sustain a 25 Hz reconstruction rate at the Tier-0; by distributing the reconstructed data to six Tier-1 Regional Centers (FNAL in US, FZK in Germany, Lyon in France, CNAF in Italy, PIC in Spain, RAL in UK) and handling catalogue issues; by redistributing data to Tier-2 centers for analysis. Simulated events, up to the digitization step, were produced prior to the DC as input for the reconstruction in the Pre-Challenge Production (PCP04).

In this paper, the model of the Tier-0 implementation used in DC04 is described, as well as the experience gained in using the newly developed data distribution management layer, which allowed CMS to successfully direct the distribution of data from Tier-0 to Tier-1 sites by loosely integrating a number of available Grid components. While developing and testing this system, CMS explored the overall functionality and limits of each component, in any of the different implementations which were deployed within DC04.

The role of Tier-1's is presented and discussed, from the import of reconstructed data from Tier-0, to the archiving on to the local mass storage system and the data distribution management to Tier-2's for analysis. Participating Tier-1's differed in available resources, set-up and configuration: a critical evaluation of the results and performances achieved adopting different strategies in the organization and management of each Tier-1 center to support CMS DC04 is presented.

Plenary / 373

Run II computing

A. Boehnlein¹¹ *FERMI NATIONAL ACCELERATOR LABORATORY***Auteur(s) contact:** cope@fnal.gov

In support of the Tevatron physics program, the Run II experiments have developed computing models and hardware facilities to support data sets at the petabyte scale, currently corresponding to 500 pb⁻¹ of data and over 2 years of production operations. The systems are complete from online data collection to user analysis, and make extensive use of central services and common solutions developed with the FNAL CD and experiment collaborating institutions, and make use of global facilities to meet the computing needs. We describe the similarities and differences between computing on CDF and D0 while describing solutions for database and database servers, data handling, movement and storage and job submission mechanisms. The facilities for production computing and analysis and the use of commodity file servers will also be described. Much of the knowledge gained from providing computing at this scale can be abstracted and applied to design and planning for future experiments with large scale computing.

Poster Session 2 - Tableau: 60 / 481

SAMGrid Experiences with the Condor Technology in Run II Computing

A. Baranovski¹ ; A. Kreymer² ; A. Kumar¹ ; A. Lyon² ; A. Nishandar³ ; G. Garzoglio² ; I. Terekhov² ; J. Trumbo¹ ; L. Loebel Carpenter¹ ; L. Lueking² ; R. Herber² ; R. Illingworth² ; R. Kennedy² ; R. St.Denis⁴ ; S. Jain³ ; S. Veseli¹ ; S. White¹ ; W. Merritt²

¹ *Fermilab*² *FERMI NATIONAL ACCELERATOR LABORATORY*³ *University of Texas, Arlington*⁴ *Glasgow University***Auteur(s) contact:** terekhov@fnal.gov

SAMGrid is a globally distributed system for data handling and job management, developed at Fermilab for the D0 and CDF experiments in Run II. The Condor system is being developed at the University of Wisconsin for management of distributed resources, computational and otherwise. We briefly review the SAMGrid architecture and its interaction with Condor, which was presented earlier. We then present our experiences using the system in production, which have two distinct aspects.

At the global level, we deployed Condor-G, the Grid-extended Condor, for the resource brokering and global scheduling of our jobs. At the heart of the system is Condor's Matchmaking Service. As a more recent work at the computing element level, we have been benefitting from the large computing cluster at the University of Wisconsin campus. The architecture of the computing facility and the philosophy of Condor's resource management have prompted us to improve the application infrastructure for D0 and CDF, in aspects such as parting with the shared file system or reliance on resources being dedicated. As a result, we have increased productivity and made our applications more portable and Grid-ready. We include some

statistics gathered from our experience. Our fruitful collaboration with the Condor team has been made possible by the Particle Physics Data Grid.

Distributed Computing Services / 460

SAMGrid Integration of SRMs

A. Baranovski¹; A. Kreymer²; A. Kumar³; A. Lyon²; A. Sill⁴; F. Ratnikov⁵; G. Garzoglio²; I. Terekhov³; J. Trumbo³; L. Loebel Carpenter³; L. Lueking²; M. Burgon-Lyon⁶; M. Leslie⁷; R. Illingworth²; R. Kennedy²; R. St.Denis⁶; S. Belforte⁸; S. Stonjek⁷; S. Veseli³; S. White³; U. Kerzel⁹; V. Bartsch⁷; W. Merritt²

¹ *Fermi National Accelerator Laboratory*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

³ *Fermilab*

⁴ *Texas Tech. University*

⁵ *Rutgers University*

⁶ *Glasgow University*

⁷ *Oxford University*

⁸ *INFN, Trieste*

⁹ *Karlsruhe University*

Auteur(s) contact: kennedy@fnal.gov

SAMGrid is the shared data handling framework of the two large Fermilab Run II collider experiments: DZero and CDF. In production since 1999 at D0, and since mid-2004 at CDF, the SAMGrid framework has been adapted over time to accommodate a variety of storage solutions and configurations, as well as the differing data processing models of these two experiments. This has been very successful for both experiments. Backed by primary data repositories of approximately 1 PB in size for each experiment, the SAMGrid framework delivers over 100 TB/day to DZero and CDF analyses at Fermilab and around the world. Each of the storage systems used with SAMGrid, however, has distinct interfaces, protocols, and behaviors. This led to different levels of integration of the various storage devices into the framework, which complicated the exploitation of their functionality and limited in some cases SAMGrid expansion across the experiments' Grid.

In an effort to simplify the SAMGrid storage interfaces, SAMGrid has

adopted the Storage Resource Manager (SRM) concept as the universal interface to all storage devices. This has simplified the SAMGrid framework, especially the implementation of storage device interactions. It prepares the SAMGrid framework for future storage solutions equipped with SRM interfaces, without the need for long and risky software integration projects. In principle, any storage device with an SRM interface can be used now with the SAMGrid framework. The integration of SRMs is an important further step towards evolving the SAMGrid framework into a co-operating collection of distinct, modular grid-oriented services. To date, SRMs for Enstore, dCache, local caches, and permanent disk locations are tested and in production use. This report outlines how the SRMs were integrated into the existing SAMGrid framework without disturbing on-going operations, and describes our operational experience with SAMGrid and SRMs in the field.

Poster Session 2 - Tableau: 38 / 451

SAMGrid Monitoring Service and its Integration with MonALisa

A. Baranovski¹ ; A. Kreymer¹ ; A. Lyon¹ ; A. Sill² ; F. Ratnikov³ ; G. Garzoglio¹ ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter⁴ ; L. Lueking¹ ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber¹ ; R. Illingworth¹ ; R. Kennedy¹ ; R. St.Denis⁵ ; S. Belforte⁷ ; S. Stonjek⁸ ; S. Veseli¹ ; S. White¹ ; U. Kerzel⁹ ; V. Bartsch⁶ ; W. Merritt¹

¹ *FERMI NATIONAL ACCELERATOR LABORATORY*

² *Texas Tech University*

³ *Rutgers University*

⁴ *Fermilab*

⁵ *Glasgow University*

⁶ *Oxford University*

⁷ *INFN/Trieste*

⁸ *FERMI NATIONAL ACCELERATOR LABORATORY/Oxford University*

⁹ *Karlsruhe University*

Auteur(s) contact: lyon@fnal.gov

The SAMGrid team is in the process of implementing a monitoring and information service, which fulfills several important roles in the operation of the SAMGrid system, and will replace the first generation of monitoring tools in the current deployments. The first generation tools are in general based on text logfiles and represent solutions which are not scalable or maintainable. The roles of the monitoring and information service are: 1) providing diagnostics for troubleshooting the operation of SAMGrid services; 2) providing support for monitoring at the level of user jobs; 3) providing runtime support for local configuration and other information currently which currently must be stored centrally (thus moving the system toward greater autonomy for the SAM station services, which include cache management and job management services); 4) providing intelligent collection of statistics in order to enable performance monitoring and tuning. The architecture of this service is quite flexible, permitting input from any instrumented SAM application or service. It will allow multiple backend storage for archiving of (possibly) filtered monitoring events, as well as real time information displays and active notification service for alarm conditions. This service will be able to export, in a configurable manner, information to higher level Grid monitoring services, such as MonALisa. We describe our experience to date with using a prototype version together with MonALisa.

Distributed Computing Services / 321

SPHINX: A Scheduling Middleware for Data Intensive Applications on a Grid

J.U. In¹ ; L. Chitnis² ; M. Kulkarni³ ; P. Avery³ ; R. Cavanaugh³ ; S. Ranka³

¹ *UNIVERSITY OF FLORIDA*

² *University of Florida*

³ *UNIVERSITY OF FLORIDA*

Auteur(s) contact: richard.cavanaugh@cern.ch

A grid consists of high-end computational, storage, and network resources that, while known a priori, are dynamic with respect to activity and availability. Efficient co-scheduling of requests to use grid resources must adapt to this dynamic environment while meeting administrative policies. We discuss the necessary requirements of such a scheduler and introduce a distributed framework called SPHINX that schedules complex, data intensive High Energy Physics and Data Mining applications in a grid environment, respecting local and global policies along with a specified level of quality of service. The SPHINX design allows for a number of functional modules and/or distributed services to flexibly schedule workflows representing multiple applications on grids. We present experimental results for SPHINX that effectively utilize existing grid middleware such as monitoring and workflow management/execution systems. These results demonstrate that SPHINX can successfully schedule work across a large number of grid sites that are owned by multiple units in a virtual organization.

Computer Fabrics / 216

SRB system at Belle/KEK

A Manabe¹; H Kuraishi²; I. Adachi³; K Ishikawa⁴; M. Mei⁵; N. Katayama¹; S Honma²; S Kawabata¹; S Yashiro¹; S. Suzuki¹; T Nakajima²; T. Sasaki³; Y Watase¹; Y. Iida³

¹ KEK

² Fujitsu

³ HIGH ENERGY ACCELERATOR RESEARCH ORGANIZATION

⁴ ISE, Chiba

⁵ IHEP, Beijing

The Belle experiment has accumulated an integrated luminosity of more than 240fb-1 so far, and a daily logged luminosity now exceeds 800pb-1. These numbers correspond to more than 1PB of raw and processed data stored on tape and an accumulation of the raw data at the rate of 1TB/day. The processed, compactified data, together with Monte Carlo simulation data for the final physics analyses amounts to more than 100TB. The Belle collaboration consists of more than 55 institutes in 14 countries and at most of the collaborating institutions, active physics data analysis programs are being undertaken. To meet these storage and data distribution demands, we have tried to adopt a resource broker, SRB. We have installed the SRB system at KEK, Australia, and other collaborating institutions and have started to share data. In this talk, experiences with the SRB system will be discussed and the performance of the system when used for data processing and physics analysis of the Belle experiment will be demonstrated.

Poster Session 2 - Tableau: 55 / 358

SRM AND GFAL TESTING FOR LCG2

E. Slabospitskaya¹; M. Sapunov¹; V. Petoukhov¹

¹ Institute for High Energy Physics, Protvino, Russia

Auteur(s) contact: elena.slabospitskaya@ihep.ru

Storage Resource Manager (SRM) and Grid File Access Library (GFAL) are GRID middleware components used for transparent access to Storage Elements. SRM provides a common interface (WEB service) to backend systems giving dynamic space allocation and file management. GFAL provides a mechanism whereby an application software can access

a file at a site without having to know which transport mechanism to use or at which site it is running.

Two separate Test Suites have been developed for testing of SRM interface v 1.1 and testing against the GFAL file system. Test Suites are written in C and Perl languages. SRM test suite: a script in Perl generates files and their replicas. These files are copied to the local SE and registered (published). Replicas of files are made to the specified SRM site. All replicas are used by the C-program. The SRM functions, such as get, put, pin, unPin etc. are tested using a program written in C. As SRMs do not perform file movement operations, the C-program transfers files using "globus-url-copy". It then compares the data files before and after transfer. GFAL test suite: as GFAL allows users to access a file in a Storage Element directly (read and write) without copying it locally, a C-program tests the implementation of POSIX I/O functions such as open/seek/read/write. A Perl script executes almost all Unix based commands: dd, cat, cp, mkdir and so on. Also the Perl script launches a stress test, creating many small files (~5000), nested directories and huge files. The investigation of interactions between the Replica Manager, the SRM and the file access mechanism will help making the Data Management software better.

Computer Fabrics / 207

ScotGrid: A prototype Tier 2 centre

A. Earl¹ ; P. Clark² ; S. Thorn^{None}

¹ UNIVERSITY OF EDINBURGH

² University of Edinburgh

Auteur(s) contact: aearl@ed.ac.uk

ScotGrid is a prototype regional computing centre formed as a collaboration between the universities of Durham, Edinburgh and Glasgow as part of the UK's national particle physics grid, GridPP. We outline the resources available at the three core sites and our optimisation efforts for our user communities. We discuss the work which has been conducted in extending the centre to embrace new projects both from particle physics and new user communities and explain our methodology for doing this.

Grid Security / 189

Secure Grid Data Management Technologies in ATLAS

A. Vaniachine¹ ; D. Malon¹ ; M. Branco²

¹ ANL

² CERN

Auteur(s) contact: miguel.branco@cern.ch

In a resource-sharing environment on the grid both grid users and grid production managers call for security and data protection from unauthorized access. To secure data management several novel grid technologies were introduced in ATLAS data management. Our presentation will review new grid technologies introduced in HEP production environment for database access through the Grid Security Infrastructure (GSI): secure GSI channel mechanisms for database services delivery for reconstruction on grid clusters behind closed firewalls; grid certificate authorization technologies for production database access control and scalable locking

technologies for the chaotic ‘on-demand’ production mode. We address the separation of file transfer process from the file catalog interaction process (file location registration, file metadata querying, etc.), database transactions capturing data integrity and the high availability fault-tolerant database solutions for the core data management tasks. We discuss the complementarities of the security model for the online and the offline computing environments; best practices (and realities) of the database users’ roles: administrators, developers, data writers, data replicators and data readers, need for elimination of the clear-text passwords; stateless and stateful protocols for the binary data transfers over secure grid data transport channels in heterogeneous grids. We present the security policies and technologies integrated in the ATLAS Production Data Management System - Don Quijote (GSI-enabled services oriented architecture, GSI proxy certificate delegation) and approaches for seamless integration of Don Quijote with POOL event collections and tag databases - while making the system non-intrusive to end-users.

Event Processing / 386

Self-Filling Histograms: An object-oriented analysis framework

B. List¹ ; J. LIST²

¹ *ETH Zurich*

² *University of Wuppertal*

Auteur(s) contact: jenny.boehme@desy.de

Analyses in high-energy physics often involve the filling of large amounts of histograms from n-tuple like data structures, e.g. RooT trees. Even when using an object-oriented framework like RooT, a the user code often follows a functional programming approach, where booking, application of cuts, calculation of weights and histogrammed quantities and finally the filling of the histogram is performed separately in different places of the program.

We will present a set of RooT based histogram classes that allow to define the histogrammed quantity, its weight and the cuts to be applied at the time of booking.

We use lightweight function object classes to define plotted quantities and cut conditions; the “self-filling” histograms hold references to these objects, and evaluate them in a fill method that thus needs no parameters. The use of function objects rather than strings to define plotted quantities and cuts permits error detection at compile rather than run time, and allows the implementation of caching mechanisms if costly computations are to be performed. Arithmetic and logical expressions are implemented by operator overloading. Histograms can be grouped in collections. We apply the visitor pattern to perform operations like filling, writing, fitting or attribute setting on such a group, without having to extend the collection class each time a new functionality is needed.

Although developed within the object oriented analysis framework of the H1 experiment, this toolkit can be used on any RooT tree.

Poster Session 1 - Tableau: 9 / 151

Simplified deployment of an EDG/LCG cluster via LCFG-UML

A. Garcia¹ ; G. Tsouloupas² ; M. Hardt³

¹ *KARLSRUHE RESEARCH CENTER (FZK)*

² *UNIVERSITY OF CYPRUS (UCY)*

³ *FORSCHUNGSZENTRUM KARLSRUHE (FZK)*

Auteur(s) contact: garcia@iwr.fzk.de, hardt@kit.edu

The clusters using DataGrid middleware are usually installed and managed by means of an “LCFG” server. Originally developed by the Univ. of Edinburgh and extended by DataGrid, this is a complex piece of software. It allows for automated installation and configuration of a complete grid site. However, installation of the “LCFG”-Server takes most of the time, thus hinder widespread use.

Our approach was to set up and preconfigure the LCFG-server inside a “User Mode Linux” (UML) instance in order to make deployment faster. The result is the “UML-LCFG-Sserver”. It is provided as a prebuilt root-filesystem image which can be up and running within only with few configuration steps. Detailed instructions and experience are also provided on the basis of tests within the CrossGrid project. Altogether UML-LCFG makes it easier for a new site to join an EDG/LCG based Grid by bypassing most of the LCFG server installation.

Poster Session 3 - Tableau: 60 / 111

Simulation and reconstruction of heavy ion collisions in the ATLAS detector.

P. Nevski¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Auteur(s) contact: nevski@bnl.gov

The ATLAS detector is a sophisticated multi-purpose detector with over 10 million electronics channels designed to study high-pT physics at LHC. Due to their high multiplicity, reaching almost hundred thousand particles per event, heavy ion collisions pose a formidable computational challenge. A set of tools have been created to realistically simulate and fully reconstruct the most difficult case of central Pb-Pb collisions (impact parameter < 1 fm) in the ATLAS detector.

A number of issues concerning extensive memory management, CPU versus memory optimization, tradeoff between data volume and physics analysis capacity have been formulated and solved. As a result we are able to predict and optimise the physics performance of the experiment and its sub-systems. We will describe the optimal dataflow organization and solutions which allowed flexible system tuning during the massive simulated data production and the analysis of tens of thousand of multi-megabyte events.

Online Computing / 437

Simulations and Prototyping of the LHCb L1 and HLT Triggers

A. Barczyk¹ ; A. SMOKER² ; B. Jost¹ ; J-P. Dufey¹ ; J. ANDERSON² ; N. Neufeld¹ ; R. McNulty³ ; T. Kechadi⁴ ; T. Shears⁵

¹ CERN

² University College Dublin, Ireland

³ UNIVERSITY COLLEGE DUBLIN

⁴ University College Dublin

⁵ University of Liverpool

Auteur(s) contact: ronan.mcnulty@cern.ch

The Level 1 and High Level triggers for the LHCb experiment are software triggers which will be implemented on a farm of about 1800 CPUs, connected to the detector read-out system by a large Gigabit Ethernet LAN with a capacity of 8 Gigabyte/s and some 500 Gigabit Ethernet links. The architecture of the readout network must be designed to maximise data throughput, control data flow, allow load balancing between the nodes and be proven to perform at scale. Issues of stability, robustness and fault tolerance are vital to the effective operation of the trigger. We report on the development and results of two independent software simulations which allow us to evaluate the performance of various network configurations and to specify the switch parameters. In order to validate the results of the simulation and to experimentally test the performance of the readout network in conditions similar to those expected at the LHC, we have constructed a hardware prototype of the LHCb Level 1 and High Level triggers. This prototype allows a scaled evaluation of our design, soak-testing, and an evaluation of the overall system response to the deliberate introduction of faults. The performance of this test-bed is described and the results compared to simulation.

Poster Session 3 - Tableau: 56 / 416

Software Management in the HARP experiment

E. Tcherniaev¹

¹ CERN

Auteur(s) contact: evgueni.tcherniaev@cern.ch

This paper discusses some key points in the organization of the HARP software. In particular it describes the configuration of the packages, data and code management, testing and release procedures.

Development of the HARP software is based on incremental releases with strict respect of the design structure.

This poses serious challenges to the software management, which has gone through essential evolution during the life of the experiment.

A progressively better understanding of the organizational issues, like the environment settings, package versioning, release procedures, etc., was achieved.

Mastering of the CVS and CMT tools, plus an essential reduction of the manual work allowed to reach the situation where one software development iteration (compilation from scratch, full testing, installation in the official area) takes only a few hours.

Distributed Computing Services / 162

Software agents in data and workflow management

A. Afaq¹ ; A. Fanfani² ; C. Charlot³ ; C. Grandi² ; D. Bonacorsi⁴ ; D. Newbold⁵ ; I. Fisk¹ ; J. Hernandez⁶ ; J. Rehn⁷ ; K. Rabbertz⁷ ; L. Tuura⁸ ; M. Ernst¹ ; O. Maroney⁵ ; P. GARCIA-ABIA⁹ ; S. Metson⁵ ; T. Barrass¹⁰ ; T. Wildish¹¹ ; W. Jank¹² ; Y. Wu¹

¹ *FNAL*

² *INFN-Bologna*

³ *IN2P3, LLR*

⁴ *CNAF-INFN*

⁵ *University of Bristol*

⁶ *CIEMAT Madrid*

⁷ *University of Karlsruhe*

⁸ *Northeastern University*

⁹ *CIEMAT*

¹⁰ *CMS, UNIVERSITY OF BRISTOL*

¹¹ *Princeton*

¹² *CERN*

CMS currently uses a number of tools to transfer data which, taken together, form the basis of a heterogeneous datagrid. The range of tools used, and the directed, rather than optimised nature of CMS recent large scale data challenge required the creation of a simple infrastructure that allowed a range of tools to operate in a complementary way.

The system created comprises a hierarchy of simple processes (named agents) that propagate files through a number of transfer states. File locations and some application metadata were stored in POOL file catalogues, with LCG LRC or MySQL backends. Agents were assigned limited responsibilities, and were restricted to communicating state in a well-defined, indirect fashion through a central transfer management database. In this way, the task of distributing data was easily divided between different groups for implementation.

The prototype system was developed rapidly, and achieved the required sustained transfer rate of ~10 MBps, with $O(10^6)$ files distributed to 6 sites from CERN. Experience with the system during the data challenge raised issues with underlying technology (MSS write/read, stability of the LRC, maintenance of file catalogues, synchronisation of filesystems) which have been successfully identified and handled. The development of this prototype infrastructure allows us to plan the evolution of backbone CMS data distribution from a simple hierarchy to a more autonomous, scalable model drawing on emerging agent and grid technology.

Core Software / 241

Software management infrastructure in the LCG Application Area

A. Pfeiffer¹ ; E. Poinsignon² ; S. Ashby² ; Y. Patois²

¹ CERN, PH/SFT

² CERN

Auteur(s) contact: andreas.pfeiffer@cern.ch

In the context of the SPI project in the LCG Application Area, a centralized s/w management infrastructure has been deployed. It comprises of a suite of scripts handling the building and validating of the releases of the various projects as well as providing a customized packaging of the released s/w. Emphasis was put on the flexibility of the packaging and distribution solution as it should cover a broad range of use-cases and needs, ranging from full packages for developers in the projects and experiments to a minimal set of libraries and binaries for specific applications running, e.g., on grid nodes. In addition, regular reviews of the QA analysis of the releases of the projects are performed and fed back to the project leaders to improve the overall quality of the software produced. The present status and future perspectives of this activity will be presented and we will show examples of quality improvement in the projects.

Poster Session 3 - Tableau: 23 / 197

Specifying Selection Criteria using C++ Expression Templates

C. Jones¹

¹ CORNELL UNIVERSITY

Generic programming as exemplified by the C++ standard library makes use of functions or function objects (objects that accept function syntax) to specialize generic algorithms for particular uses. Such separation improves code reuse without sacrificing efficiency. We employed this same technique in our combinatoric engine: DChain. In DChain, physicists combine lists of child particles to form a list of parent hypotheses. E.g., `d0 = pi.plus() * K.minus()`. The selection criteria for the hypothesis is defined in a function or function object that is passed to the list's constructor.

However, C++ requires that functions and class declarations be defined outside the scope of a function.

Therefore physicists are forced to separate the code that defines the combinatorics from the code that sets the selection criteria. We will discuss a technique using C++ expression templates to allow users to define function objects using a mathematical expression directly in their main function, e.g.,

```
func = (sqrt( beamEnergy*beamEnergy - vPMagvPMag) >= 5.1*k_GeV).
```

Use of such techniques can greatly decrease the coding 'excess' needed to perform an analysis.

Poster Session 3 - Tableau: 42 / 240

Status and Plans of the LCG PI Project

A. Pfeiffer¹ ; H.C. Lee² ; L. Moneta¹ ; V. Innocente¹ ; W. Ueng²

¹ CERN, PH/SFT

² Academia Sinica, Taipei

Auteur(s) contact: andreas.pfeiffer@cern.ch

In the context of the LHC Computing Grid (LCG) project, the Applications Area develops and maintains that part of the physics applications software and associated infrastructure that is shared among the LHC experiments.

The Physicist Interface (PI) project of the LCG Application Area encompasses the interfaces and tools by which physicists will directly use the software. In collaboration with users from the experiments, work has concentrated on the Analysis Services subsystem, where implementations of the AIDA interfaces for (binned and unbinned) histogramming, fitting and minimization as well as manipulation of tuples have been developed and adapted. In addition, bindings of these interfaces to the Python interpreted language have been done using the dictionary subsystem of the SEAL project.

The actual status and the future planning of the project will be presented.

Poster Session 1 - Tableau: 10 / 410

Status of the alignment calibrations in the ATLAS-Muon experiment

B. DUBOUE¹ ; C. GUYOT¹ ; J.C BARRIERE¹ ; J.P. SCHULLER¹ ; O. CLOUE¹ ; P. PERRIN¹ ; P. PONSOT¹ ; P. SCHUNE¹ ; V. GAUTARD² ; Y. REINERT¹

¹ DAPNIA, CEA-Saclay, France

² CEA-SACLAY

Auteur(s) contact: gautard@cea.fr

ATLAS is a particle detector which will be built at CERN in Geneva. The muon detection system is made up among other things, of 600 chambers measuring 2 to 6 m² and 30 cm thick. The chambers' position must be known with an accuracy of +/- 30 μ m for translations and +/-100 μ m for rotations for a range of +/- 5mm and +/-5mrad. In order to fulfill these requirements, we have designed different optical sensors.

Due to (i) the very high accuracy required, (ii) the number of sensors (over 1000) and (iii) the different type of sensors, we developed one user interface which manages among other things several control command software. Each of this software is associated with an accurate calibration bench. In this conference, we will present only the most complex one which combines command control, an analysis module, real time processing and database access. These softwares are now currently used for sensors calibration.

Computer Fabrics / 461

StoRM: grid middleware for disk resource management

A. Ghiselli¹ ; F. Donno² ; L. Magnoni¹ ; R. Zappi¹

¹ *INFN-CNAF*

² *INFN*

Auteur(s) contact: riccardo.zappi@cnaif.infn.it

Within a Grid the possibility of managing storage space is fundamental, in particular, before and during application execution. On the other hand, the increasing availability of highly performant computing resources raises the need for fast and efficient I/O operations and drives the development of parallel distributed file systems able to satisfy these needs granting access to distributed storage.

The demand of POSIX compliant access to storage and the need to have a uniform interface for both Grid integrated and pure vanilla applications stimulate developers to investigate the possibility to integrate already existing filesystems into a Grid infrastructure, allowing users to take advantage of storage resources without being forced to change their applications.

This paper describes the design and implementation of StoRM, a storage resource manager (SRM) for disk only. Through StoRM an application can reserve and manage space on disk storage systems. It can then access the space either in a Grid environment or locally in a transparent way via classic POSIX calls.

The StoRM architecture is based on a pluggable model in order to easily add new functionalities. The StoRM implementation uses now filesystems such as GPFS or LUSTRE. The StoRM prototype includes space reservation functionalities that complement SRM space reservation to allow applications to directly access/use the managed space through POSIX calls. Moreover, StoRM includes quota management and a space guard. StoRM will serve as policy enforcement point (PEP) for the Grid Policy Management System over disk resources. The experimental results obtained are promising.

Computer Fabrics / 107

Storage Resource Manager

D. Petravick¹ ; J. Bakken¹ ; T. Perelmutov¹

¹ *FERMI NATIONAL ACCELERATOR LABORATORY*

Auteur(s) contact: timur@fnal.gov

Storage Resource Managers (SRMs) are middleware components whose function is to provide dynamic space allocation and file management on shared storage components on the Grid. SRMs support protocol negotiation and reliable replication mechanism. The SRM standard allows independent institutions to implement their own SRMs, thus allowing for a uniform access to heterogeneous storage elements. SRMs leave the policy decision to be made independently by each implementation at each site. Resource Reservations made through SRMs have limited lifetimes and allow for automatic collection of unused resources thus preventing clogging of storage systems with “forgotten” files.

The storage systems can be classified on basis of their longevity and persistence of their data. Data can also be temporary or permanent. To support these notions, SRM defines Volatile, Durable and Permanent types of files and spaces. Volatile files can be removed by the system to make space for new files upon the expiration of their lifetimes. Permanent files are expected to exist in the storage system for the lifetime of the storage system. Finally Durable files have both the lifetime associated with them and a mechanism of notification of owners and administrators of lifetime expiration, but cannot be deleted automatically by the system and require explicit removal.

Fermilab’s data handling system uses the SRM management interface, the dCache Distributed Disk Cache and the Enstore Tape Storage System as key components to satisfy current and future user requests.

Storage Resource Manager specification is a result of international collaborative effort by representatives of JLAB, LBNL, FNAL, EDG-WP2 and EDG-WP5.

Distributed Computing Services / 345

Storage Resource Managers at Brookhaven

O. Rind¹ ; R. Popescu¹ ; S. O'Hare¹ ; Z. Liu¹

¹ *BROOKHAVEN NATIONAL LABORATORY*

Auteur(s) contact: rind@bnl.gov

Providing Grid applications with effective access to large volumes of data residing on a multitude of storage systems with very different characteristics prompted the introduction of storage resource managers (SRM). Their purpose is to provide consistent and efficient wide-area access to storage resources unconstrained by their particular implementation (tape, large disk arrays, dispersed small disks). To assess their viability in the context of the US Atlas Tier 1 facility at Brookhaven, two implementations of SRM were tested: dCache (FNAL/DESY joint project) and HRM/DRM (NERSC Berkeley). Both systems included a connection to the local HPSS mass data store providing Grid access to the main tape repository. In addition, dCache offered storage aggregation of dispersed small disks (local drives on computing farm nodes). An overview of our experience with both systems will be presented, including details about configurations, performance, inter-site transfers, interoperability and limitations.

Plenary / 516

Summary - Computer Fabrics

T. SMITH^{None}

Plenary / 513

Summary - Core Software

P. CANAL^{None}

Plenary / 514

Summary - Distributed Computing Services

M. Lamanna¹

¹ *CERN*

Plenary / 515

Summary - Distributed Computing Systems and Experiences

D. Olson¹

¹ *Lawrence Berkeley Laboratory*

Plenary / 512

Summary - Event Processing

S. GOWDY^{None}

Plenary / 511

Summary - Online Computing

P. VANDE VYVRE^{None}

Plenary / 517

Summary - Wide Area Networking

P. CLARKE^{None}

Distributed Computing Systems and Experiences / 287

Super scaling PROOF to very large clusters

F. Rademakers¹ ; G. Roland² ; K. Gulbrandsen² ; M. Ballintijn² ; P. Canal³ ; R. Brun¹

¹ *CERN*

² *MIT*

³ *Fermilab*

Auteur(s) contact: maartenb@mit.edu

The Parallel ROOT Facility, PROOF, enables a physicist to analyze and understand very large data sets on an interactive time scale. It makes use of the inherent parallelism in event data and implements an architecture that optimizes I/O and CPU utilization in heterogeneous clusters with distributed storage. Scaling to many hundreds of servers is essential to process tens or hundreds of gigabytes of data interactively. This is supported by the industry trend to pack more CPU's into single systems and to create bigger clusters by increasing the number of systems per rack. We will describe the latest developments in PROOF and the development of a standardized benchmark for PROOF clusters. The benchmark is self contained and measures the network, the I/O and the processing characteristics of

a cluster. We will present the comprehensive results of the benchmark for several clusters, demonstrating the performance and scalability of PROOF on very large clusters.

Core Software / 367

Supporting the Development Process of the DataGrid Workload Management System Software with GNU autotools, CVS and RPM

Auteur(s): E. Ronchieri¹ ; F. Prelz² ; R. Peluso³

Co-auteur(s) A. Gianelle³ ; A. Guarise⁴ ; A. Krenek⁵ ; A. Maraschini⁶ ; A. Werbrouck⁴ ; B. Cantalupo⁷ ; D. Colling⁸ ; D. Kouril⁵ ; F. Giacomini¹ ; F. Pacini⁶ ; G. Avellino⁶ ; J. Pospisil⁵ ; J. Sitera⁵ ; L. Matyska⁵ ; M. Mezzadri² ; M. Mulac⁵ ; M. Pappalardo⁹ ; M. Ruda⁵ ; M. Sgaravatto³ ; M. Vocu⁵ ; R. Piro⁴ ; S. Beco⁷ ; S. Monforte⁹ ; Z. Kabelac⁵ ; Z. Salvat⁵

¹ INFN CNAF

² INFN Milano

³ INFN Padova

⁴ INFN Torino

⁵ CESNET

⁶ Datamat Ingegneria dei Sistemi S.p.A

⁷ Datamat Ingegneria dei Sistemi S.p.A.

⁸ Imperial College, London

⁹ INFN Catania

Auteur(s) contact: elisabetta.ronchieri@cnaif.infn.it

We described the process for handling software builds and releases for the Workload Management package of the DataGrid project. The software development in the project was shared among nine contractual partners, in seven different countries, and was organized in work-packages covering different areas.

In this paper, we discuss how a combination of Concurrent Version System, GNU autotools and other tools and practices was organised to allow the development, build, test and distribution of the DataGrid Workload Management System. This is not only characterised by a rather high internal geographic and administrative dispersion (four institutions with developers at nine different locations in three countries), but by the fact we had to integrate and interface to a dozen of third-party code packages coming from different sources, and to the software products coming from other three development work-packages internal to the project.

A high level of central co-ordination needed to be maintained for project-wide steering, and this had also to be reflected in the software development infrastructure, while maintaining ease-of-use for distributed developers and automated procedures wherever possible.

Event Processing / 315

Synergia: A Modern Tool for Accelerator Physics Simulation

J. Amundson¹ ; P. Spentzouris²

¹ *Fermi National Accelerator Laboratory*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

Auteur(s) contact: spentz@fnal.gov

Computer simulations play a crucial role in both the design and operation of particle accelerators. General tools for modeling single-particle accelerator dynamics have been in wide use for many years. Multi-particle dynamics are much more computationally demanding than single-particle dynamics, requiring supercomputers or parallel clusters of PCs. Because of this, simulations of multi-particle dynamics have been much more specialized. Although several multi-particle simulation tools are now available, they tend to cover a narrow range of topics. Most also present difficulties for the end user ranging from platform portability to arcane interfaces.

In this presentation, we discuss Synergia, a multi-particle accelerator simulation tool developed at Fermilab, funded by the DOE SciDAC program. Synergia was designed to cover a variety of physics processes while presenting a flexible and humane interface to the end user. It is a hybrid application, primarily based on the existing packages mxyzptlk/beamline and Impact. Our presentation covers Synergia's physics capabilities and human interface. We focus on the computational problems we encountered and solved in the process of building an application out of codes written in Fortran 90, C++, and wrapped with a Python front-end. We also discuss some approaches we have used in the visualization of the high-dimensional data that comes out of a particle accelerator simulations, especially our work with OpenDX.

Wide Area Networking / 51

Test of data transfer over an international network with a large RTT

H. Matsumoto¹ ; H. Sakamoto¹ ; I. Ueda¹ ; J. Tanaka¹ ; M. Ishino¹ ; S.Y. SUZUKI² ; T. Mashimo¹

¹ *ICEPP, UNIVERSITY OF TOKYO*

² *KEK*

Auteur(s) contact: junichi.tanaka@cern.ch

We have measured the performance of data transfer between CERN and our laboratory, ICEPP, at the University of Tokyo in Japan. The ICEPP will be one of the so-called regional centers for handling the data from the ATLAS experiment which will start data taking in 2007. More than petabytes of data are expected to be generated from the experiment each year. It is therefore essential to achieve a high throughput of data transfer over the long-distance network connection between CERN and ICEPP. A connection with several gigabits per second is now available between the two sites. The round trip time, however, reaches about 300 msec. Moreover the connection is not dedicated to us. Due to the large latency and other traffic on the same network, it is not easy to fully exploit the available bandwidth. We have measured the performance of the network connection using tools such as iperf, bftf, and gridftp with various TCP parameters, Linux kernel versions and so on. We have examined factors limiting the speed and tried to improve the throughput of the data transfer.

In this talk we report on the results of our measurements and investigations.

Event Processing / 110

Test of the ATLAS Inner Detector reconstruction software using combined test beam data

M.J. Costa¹ ; W. Liebig¹

¹ *CERN*

Auteur(s) contact: wolfgang.liebig@cern.ch, maria.jose.costa@cern.ch

The athena software framework for event reconstruction in ATLAS will be employed to analyse the data from the 2004 combined test beam. In this combined test beam, a slice of the ATLAS detector is operated and read out under conditions similar to future LHC running, thus providing a test-bed for the complete reconstruction chain. First results for the ATLAS InnerDetector will be presented.

In particular, the reading of the bytestream data inside athena, the monitoring tasks, the alignment techniques and all the different online and offline reconstruction algorithms will be fully tested with real data. Their performance will be studied and results compared to simulated data, which has been generated specifically for the test beam layout.

Online Computing / 266

Testbed Management for the ATLAS TDAQ

M. ZUREK¹

¹ *CERN, IFJ KRAKOW*

Auteur(s) contact: marian.zurek@cern.ch

The talk presents the experience gathered during the testbed administration (~100 PC and 15+ switches) for the ATLAS Experiment at CERN.

It covers the techniques used to resolve the HW/SW conflicts, network related problems, automatic installation and configuration of the cluster nodes as well as system/service monitoring in the heterogeneous dynamically changing cluster environment.

Techniques range from manual actions to the fully automated procedures based on tools like Kickstart, SystemImager, Nagios, MRTG and Spectrum. Booting diskless nodes using EtherBoot, PXEboot is also investigated as a possible technique of managing Atlas Production Farms.

Kernel customization techniques (building, deploying, distribution policy) allow users to freely choose proffered kernel flavors without sysadmin intervention. At the same time administrator retains full control over entire testbed.

The overall experience has shown that the proper use of the open-source tools addresses very well the needs of the ATLAS Trigger DAQ community. This approach may also be interesting for addressing certain aspects of GRID Farm Management.

Poster Session 2 - Tableau: 11 / 113

Testing the CDF Distributed Computing Framework

A. Baranovski¹ ; A. Kreymmer² ; A. Lyon² ; A. Sill³ ; F. Ratnikov⁴ ; G. Garzoglio² ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter¹ ; L. Lueking² ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber² ; R. Illingworth² ; R. Kennedy² ; R. St.Denis⁵ ; S. Belforte⁷ ; S. Stonjek⁸ ; S. Veseli¹ ; S. White¹ ; U. Kerzel⁹ ; V. Bartsch¹⁰ ; W. Merritt²

¹ *Fermilab*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

³ *Texas Tech. University*

⁴ *Rutgers University*

⁵ *Glasgow University*

⁶ *Oxford University*

⁷ *INFN, Trieste*

⁸ *FNAL*

⁹ *Karlsruhe University*

¹⁰ *OXFORD UNIVERSITY*

To distribute computing for CDF (Collider Detector at Fermilab) a system managing local compute and storage resources is needed. For this purpose CDF will use the DCAF (Decentralized CDF Analysis Farms) system which is already at Fermilab. DCAF has to work with the data handling system SAM (Sequential Access to data via Metadata). However, both DCAF and SAM are mature systems which have not yet been used in combination, and on top of this DCAF has only been installed at Fermilab and not on local sites. Therefore tests of the systems are necessary to test the interplay of the data handling with the farms, the behaviour of the off-site DCAFs and the user friendliness of the whole system. The tests are focussed on the main tasks of the DCAFs, like Monte Carlo generation and stores, as well as the readout of data files and connected data handling. To achieve user friendliness the SAM station environment has to be common to all stations and adaptations to the environment have to be made.

Distributed Computing Systems and Experiences / 506

The ALICE Data Challenge 2004 and the ALICE distributed analysis prototype

A. Morsch¹ ; A. Peters² ; F. Carminati¹ ; F. Rademakers¹ ; L. Betev¹ ; P. Buncic¹ ; P. Hristov¹

¹ *CERN*

² *ce*

Auteur(s) contact: andreas.joachim.peters@cern.ch, federico.carminati@cern.ch

During the first half of 2004 the ALICE experiment has performed a large distributed computing exercise with two major objectives: to test the ALICE computing model, included distributed analysis, and to provide data sample for a refinement of the

ALICE Jet physics Monte-Carlo studies. Simulation reconstruction and analysis of several hundred thousand events were performed, using the heterogeneous resources of tens of computer centres worldwide. These resources belong to different GRID systems and were steered by the AliEn (ALICE Environment) framework, acting as a meta-GRID. This has been a very thorough test of the middleware of AliEn and LCG (LCG-2 and grid.it resources) and their compatibility. During the Data Challenge more than 1,500 jobs run in parallel for several weeks. More than 50 TB of data have been produced and analysed worldwide in one of the major exercises of this kind run to date. ALICE has developed an analysis system based on AliEn and ROOT. This system starts with a metadata selection in the AliEn file catalogue, followed by a computation phase. Analysis jobs are sent where the data is, thus minimising data movement. The control is performed by an intelligent workload management system. The analysis can be done either via batch or interactive jobs. The latter are “spawned” on remote systems and report the results back to the user workstation. The talk will describe the ALICE experience with this large-scale use of the Grid, the major lessons learned and the consequences for the ALICE computing model.

Online Computing / 62

The ALICE Experiment Control System

A. VASCOTTO¹ ; C. SOOS¹ ; F. Carena² ; J-C. MARIN¹ ; K. SCHOSSMAIER¹ ; P. VANDE VYVRE¹ ; R. DIVIA¹ ; S. CHAPELAND¹ ; W. CARENA¹

¹ CERN, Geneva, Switzerland

² CERN

Auteur(s) contact: franco.carena@cern.ch

The Experiment Control System (ECS) is the top level of control of the ALICE experiment.

Running an experiment implies performing a set of activities on the online systems that control the operation of the detectors. In ALICE, online systems are the Trigger, the Detector Control Systems (DCS), the Data-Acquisition System (DAQ) and the High-Level Trigger (HLT).

The ECS provides a framework in which the operator can have a unified view of all the online systems and perform operations on the experiment seen as a set of detectors.

ALICE has adopted a hierarchical -yet loose- architecture, in which the ECS is a layer sitting above the online systems, still preserving their autonomy to operate independently. The interface between the ECS and the online systems applies a powerful paradigm based on inter-communicating objects. The behavioural aspects of the ECS are described using a finite-state machine model.

The ALICE experiment must be able to run either as a whole (during the physics production) or as a set of independent detectors (for installation and commissioning). The ECS provides all the features necessary to split the experiment into partitions, containing one or more detectors, which can be operated independently and concurrently.

This paper will present the architecture of the ALICE ECS, its current status and the practical experience acquired at the test beams.

Online Computing / 434

The ALICE High Level Trigger

A. Vestbo¹ ; A. Wiebalck² ; B. Skaali³ ; C. Loizides⁴ ; D. Rohrich¹ ; H. Helstrup⁵ ; H. Tilsner² ; K. Ullaland¹ ; M. Richter¹ ; R. Stock⁴ ; T. Alt² ; T. Vik³ ; T.M. Steinbeck² ; V. Lindenstruth²

¹ *Department of Physics and Technology, University of Bergen, Norway*

² *Kirchhoff Institute for Physics, University of Heidelberg, Germany*

³ *Department of Physics, University of Oslo, Norway*

⁴ *Institute for Nuclear Physics, University of Frankfurt, Germany*

⁵ *Faculty of Engineering, Bergen University College, Norway*

Auteur(s) contact: matthias.richter@ift.uib.no

The ALICE experiment at LHC will implement a High Level Trigger System, where the information from all major detectors are combined, including the TPC, TRD, DIMUON, ITS etc. The largest computing challenge is imposed by the TPC, requiring realtime pattern recognition. The main task is to reconstruct the tracks in the TPC, and in a final stage combine the tracking information from all detectors. Based on the physics observables selective readout is done by generation of a software trigger (High Level Trigger), capable of selecting interesting (sub)events from the input data stream. Depending on the physics program various processing options are currently being developed, including region of interest processing, rejecting events based on software trigger and data compression schemes. Examples of such triggers are verification of candidates for high-pt dielectron heavy-quarkonium decays, momentum filter to enhance the open-charm signal, high-pt jets selection etc.

Technically the HLT system entails a very large scale processing farm with about 1000 active processors. The input data stream is designed for 25 GB/sec. The system nodes will be interfaced to the local data concentrators of the DAQ system via optical fibers receiving a copy of the raw data.

The optical fibers will be connected to the PCI-bus of HLT nodes using a custom PCI card. These cards provide a co-processor functionality for the first steps of the pattern recognition.

The talk will give an overview of the HLT project and will focus on the latest results regarding efficient data compression and trigger performance.

Distributed Computing Systems and Experiences / 300

The ARDA Prototypes

A. Demichev¹ ; A. Maier² ; A. Peters² ; B. Koblitz² ; D. Feichtinger³ ; D. Liko² ; F. Orellana² ; J. Andreeva² ; J. Herrala² ; J. MOSCICKI² ; J. Moscicki² ; M. Lamanna² ; T. Chen⁴ ; V. Pose⁵ ; W. Ueng⁴

¹ *Moscow State University*

² *CERN*

³ *Swiss Institute of Particle Physics*

⁴ *Academica Sinica, Nanking*

⁵ *JINR, Dubna*

Auteur(s) contact: julia.andreeva@cern.ch

The ARDA project was started in April 2004 to support the four LHC experiments (ALICE, ATLAS, CMS and LHCb) in the implementation of individual production and analysis environments based on the EGEE middleware.

The main goal of the project is to allow a fast feedback between the experiment and the middleware development teams via the construction and the usage of end-to-end prototypes allowing users to perform analyses out of the present data sets from recent montecarlo productions.

We present the status of the integration of the EGEE prototype Grid middleware into the analysis environment of the four LHC experiments. First an overview is given on the individual architectures of the four experiments' prototypes with a strong focus on how the EGEE middleware is incorporated into the framework. We outline common points in the usage of the middleware and try to point out differences in the decisions taken by the experiments on the inclusion of different parts of the EGEE software. We will conclude by presenting the first feedback from the usage of these analysis environments.

Poster Session 2 - Tableau: 5 / 161

The ATLAS Computing Model

R. JONES¹

¹ LANCAS

Auteur(s) contact: roger.jones@cern.ch

The ATLAS Computing Model is under continuous active development. Previous exercises focussed on the Tier-0/Tier-1 interactions, with an emphasis on the resource implications and only a high-level view of the data and workflow. The work presented here considerably revises the resource implications, and attempts to describe in some detail the data and control flow from the High Level Trigger farms all the way through to the physics user. The model draws from the experience of previous and running experiments, but will be tested in the ATLAS Data Challenge 2 (DC2, described in other abstracts) and in the ATLAS Combined Testbeam exercises.

An important part of the work is to devise the measurements and tests to be run during DC2. DC2 will be nearing completion in September 2004, and the first assessments of the performance of the computing model in scaled slice tests will be presented.

Poster Session 1 - Tableau: 34 / 219

The ATLAS DAQ system

G. unel¹ ; TDAQ ATLAS²

¹ UNIVERSITY OF CALIFORNIA AT IRVINE AND CERN

² ATLAS

Auteur(s) contact: gokhan.unel@cern.ch

The 40 MHz collision rate at the LHC produces ~25 interactions per bunch crossing within the ATLAS detector, resulting in terabytes of data per second to be handled by the detector electronics and the trigger and DAQ system. A Level 1 trigger system based on custom designed and built electronics will reduce the event rate to 100 kHz.

The DAQ system is responsible for the readout of the detector specific electronics via 1600 point to point links hosted by Readout Subsystems, the collection and provision of "Region of Interest data" to the Level 2 trigger, the building of events accepted by the Level 2 trigger and their subsequent input to the Event Filter system where they are subject to further selection criteria. Also the DAQ provides the functionality for the configuration, control, information exchange and monitoring of the whole ATLAS detector.

The baseline ATLAS DAQ architecture and its implementation will be introduced. In this implementation, the configuration, control, information exchange and monitoring functionalities are provided with CORBA; the control aspects are handled by an expert system based on CLIPS and the data connection between 150 Readout Subsystems, up to 500 Level 2 Processing Units and to 80 Event building nodes is done Gigabit Ethernet network technology.

The experience from using the DAQ system in a combined test beam environment where all ATLAS subdetectors are participating will be presented. The current performances of some DAQ components as measured in the laboratory environment will be summarized. Some results from the large scale functionality tests, on a system of a 300 nodes, aimed at understanding the scalability of the current implementation will also be shown.

Core Software / 424

The AliRoot framework, status and perspectives

A. Morsch¹ ; F. Carminati¹ ; F. Rademakers¹ ; K. Safarik¹ ; P. Buncic¹ ; P. Hristov¹

¹ CERN

Auteur(s) contact: federico.carminati@cern.ch

The ALICE collaboration at the LHC is developing since 1998 an OO offline framework, written entirely in C++.

In 2001 a GRID system (AliEn - ALICE Environment) has been added and successfully integrated with ROOT and the offline. The resulting combination allows ALICE to do most of the design of the detector and test the validity of its computing model by performing large scale Data Challenges, using OO technology in a distributed framework. The early migration of all ALICE users to C++ and the adoption of advanced software development techniques are two of the strong points of the ALICE offline strategy. The offline framework is heavily based on virtual interfaces, which allows the use of different generators and even different Monte-Carlo transport codes with no change in the framework or the scoring, reconstruction and analysis code. This talk presents a review of the development path, current status and future perspectives of the ALICE Offline environment.

Online Computing / 422

The Architecture of the ZEUS Second Level Global Tracking Trigger

M. Sutton¹ ; ZEUS GTT group^{None}

¹ UNIVERSITY COLLEGE LONDON

Auteur(s) contact: sutt@cern.ch

The architecture and performance of the ZEUS Global Track Trigger (GTT) are described. Data from the ZEUS silicon Micro Vertex detector's HELIX readout chips, corresponding to 200k channels, are digitized by 3 crates of ADCs and PowerPC VME board computers push cluster data for second level trigger processing and strip data for event building via Fast and GigaEthernet network connections. Additional tracking information from the central tracking chamber and forward straw tube tracker are interfaced into the 12 dual CPU PC farm of the global track trigger where track and vertex finding is performed by separately threaded algorithms. The system is data driven at the ZEUS first level trigger rates <500Hz, generating trigger results after a mean time of 10ms. The GTT integration into the ZEUS second level trigger and recent performance are reviewed.

Core Software / 108

The Athena Control Framework in Production, New Developments and Lessons Learned

C. Leggett¹ ; D. Quarrie¹ ; M. Marino¹ ; P. Calafiura¹ ; W. Lavrijsen¹

¹ LBNL

Auteur(s) contact: pcalafiura@lbl.gov

Athena is the Atlas Control Framework, based on the common Gaudi architecture, originally developed by LHCb. In 2004 two major production efforts, the Data Challenge 2 and the Combined Test-beam reconstruction and analysis were structured as Athena applications. To support the production work we have added new features to both Athena and Gaudi: an "Interval of Validity" service to manage time-varying conditions and detector data; a History service, to manage the provenance information of each event data object; and a toolkit to simulate and analyze the overlay of multiple collisions during the detector sensitive time (pile-up). To support the analysis of simulated and test-beam data in athena we have introduced a python-based scripting interface, based on the CERN LCG tools PyLCGDict, PyRoot and PyBus. The scripting interface allows to fully configure any athena component, interactively browse and modify this configuration, as well as examine the content of any data object in the event or detector store.

Event Processing / 168

The Atlantis event visualisation program for the ATLAS experiment

C. Timmermans¹ ; D. Petrusca² ; E. JANSEN¹ ; F. Crijns¹ ; G. Taylor³ ; H. Drevermann⁴ ; J. Couchman⁵ ; J. Drohan⁵ ; N. Konstantinidis⁶ ; P. Klok¹ ; Z. MAXA⁵

¹ University of Nijmegen

² Seigen/CERN

³ UC Santa Cruz

⁴ CERN

⁵ *University College London*

⁶ *UNIVERSITY COLLEGE LONDON*

Auteur(s) contact: nikolaos.konstantinidis@cern.ch

We describe the philosophy and design of Atlantis, an event visualisation program for the ATLAS experiment at CERN. Written in Java, it employs the Swing API to provide an easily configurable Graphical User Interface.

Atlantis implements a collection of intuitive, data-orientated 2D projections, which enable the user to quickly understand and visually investigate complete ATLAS events. Event data is read in from XML files produced by a dedicated algorithm running in the ATLAS software framework ATHENA, and translated into internal data objects. Within the same main canvas area, multiple views of the data can be displayed with varying size and position. Interactions such as zoom, selection and query can occur between these views using Drag and Drop.

Associations between data objects as well as the values of their member variables provide criteria upon which the Atlantis user may filter a full Atlas event. By choosing whether or not to show certain data and, if so, in what colour, a more personalised and useful display may be obtained. The user can dynamically create and manage their own associations and perform context dependent operations upon them.

Poster Session 2 - Tableau: 56 / 350

The BABAR Analysis Task Manager

Auteur(s): W. Roethel¹

Co-auteur(s) Alessandra Forti²; Antonio Ceseracciu³; David E. Hutchcroft⁴; Dmitry Bukin⁵; Dmytro Kovalskyi⁶; Douglas Smith⁷; Gregory Dubois-Felsmann⁸; Paul Stephen Jackson⁹; Tim Adye¹⁰

¹ *University of California, Irvine*

² *University of Manchester*

³ *Universita' di Padova*

⁴ *University of Liverpool*

⁵ *Budker Institute of Nuclear Physics*

⁶ *University of Maryland*

⁷ *Stanford Linear Accelerator Center*

⁸ *California Institute of Technology*

⁹ *University of London, Royal Holloway and Bedford New College*

¹⁰ *Rutherford Appleton Laboratory*

Auteur(s) contact: douglas@slac.stanford.edu, roethel@slac.stanford.edu

The new BaBar bookkeeping system comes with tools to directly support data analysis tasks. This Task Manager system acts as an interface between datasets defined in the bookkeeping system, which are used as input to analyzes, and the offline analysis framework. The Task Manager organizes the processing of the data by creating specific jobs to be either submitted to a batch system, or run in the background on a local desktop, or laptop. The current system has been designed to support pbs and lsf batch systems. Changes to defined datasets due production is directly supported by the Task Manager, where new collections that add to a dataset or replace other collections are automatically detected, allowing an analysis at any time to be

up-to-date with the latest available data. The output of tasks, whether new data collections, ntuple/hbook files, or text files, can be put back into a collections bookkeeping system or stored in the private Task Manager database. Currently MySQL and Oracle relational databases are supported. The BABAR Task Manager has been in use for data production since January this year, and the schema of the working system will be presented.

Plenary / 491

The BIRN Project: Distributed Information Infrastructure and Multi-scale Imaging of the Nervous System (BIRN = Biomedical Informatics Research Network)

M. Ellisman¹

¹ *National Center for Microscopy and Imaging Research of the Center for Research in Biological Systems - The Department of Neurosciences, University of California San Diego School of Medicine - La Jolla, California - USA*

The grand goal in neuroscience research is to understand how the interplay of structural, chemical and electrical signals in nervous tissue gives rise to behavior. Experimental advances of the past decades have given the individual neuroscientist an increasingly powerful arsenal for obtaining data, from the level of molecules to nervous systems. Scientists have begun the arduous and challenging process of adapting and assembling neuroscience data at all scales of resolution and across disciplines into computerized databases and other easily accessed sources. These databases will complement the vast structural and sequence databases created to catalogue, organize and analyze gene sequences and protein products. The general premise of the neuroscience goal is simple; namely that with “complete” knowledge of the genome and protein structures accruing rapidly we next need to assemble an infrastructure that will facilitate acquisition of an understanding for how functional complexes operate in their cell and tissue contexts. Our U.C. San Diego-based group is leading several interdisciplinary projects around this grand challenge. We are evolving a shared infrastructure that allows for mapping molecular and cellular brain anatomy in the context of a shared multi-scale mouse brain atlas system, the Cell-Centered Database (CCDB). Complementary to these neuroinformatics activities at the National Center for Microscopy and Imaging Research in San Diego (NCMIR) we have developed new molecular labeling methods compatible with advanced ultra-wide field laser-scanning light microscopy and multi-resolution 3 dimensional electron microscopy. These new labeling and imaging methods are being used to populate the CCDB, using as a driver mouse models of neurological and neuropsychiatric disorders. The informatics framework is facilitating cooperative work by distributed teams of scientists engaged in focused collaborations aimed to deliver new fundamental understanding of structures on the scale of 1 nm³ to 10⁷’s of μm³, a dimensional range that encompasses macromolecular complexes, organelles, and multi-component structures like synapses and the cellular interactions in the context of the complex organization of the entire nervous system. This is a unique and pioneering effort that links new neuroscience techniques and revolutionary advances in information technology. Database federation tools are critical to the scalability of these efforts and future development plans will be described in the context of the NIH-supported project to create a new framework for collaboration and data integration in the Biomedical Informatics Research Network (BIRN). BIRN is the leading example of a virtual database effort that is using the challenge of federating multi-scale distributed data about the nervous systems to help guide the evolution of an International Cyberinfrastructure serving all science disciplines, including biomedicine.

Poster Session 3 - Tableau: 16 / 256

The Binary Cascade

G. Folger¹ ; H.P. Wellisch¹

¹ CERN

Auteur(s) contact: gunter.folger@cern.ch, hpw@free.fr

Geant4 is a toolkit for the simulation of the passage of particles through matter. Amongst its applications are hadronic calorimeters of LHC detectors and simulation of radiation environments. For these types of simulation, a good description of secondaries generated by inelastic interactions of primary nucleons and pions is particularly important.

The Geant4 Binary Cascade is a hybrid between a classical intra-nuclear cascade and a QMD model for the simulation of inelastic scattering of pions, protons and neutrons, and light ions of intermediate energies off nuclei. The nucleus is modeled by individual nucleons bound in the nuclear potential. Binary collisions of projectiles or projectile constituents and secondaries with single nucleons, resonance production, and decay are simulated according to measured, parametrised or calculated cross sections. Pauli's exclusion principle, i.e. blocking of interactions due to Fermi statistics, reduces the free cross section to an effective intra-nuclear cross section. Secondary particles are allowed to further interact with remaining nucleons.

We will describe the modeling, and give an overview of the components of the model, their object oriented design, and implementation.

Poster Session 3 - Tableau: 64 / 79

The CEDAR Project

B. Waugh¹ ; J. Butterworth¹ ; J. Stirling² ; M. Whalley³ ; S. BUTTERWORTH⁴

¹ UCL, London

² IPPP, Durham

³ IPPP, UNIVERSITY OF DURHAM

⁴ UCL, London, UK

Auteur(s) contact: m.r.whalley@durham.ac.uk

We will describe the plans and objectives of the recently funded PPARC(UK) e-science project, the Combined E-Science Data Analysis Resource for High Energy Physics (CEDAR), which will combine the strengths of the well established and widely used HEPDATA library of HEP data and the innovative JETWEB Data/Monte Carlo comparison facility built on the HZTOOL package and which exploits developing grid technology. The current status and future plans of both of these individual sub-projects within the CEDAR framework are described showing how they will cohesively provide a) an extensive archive of Reaction Data, b) validation and tuning of Monte Carlo programmes against the Reaction Data sets, and c) a validated code repository for a wide range of HEP code such as parton distribution functions and other calculation codes used by particle physicists. Once established it is envisaged CEDAR will become an important GRID tool used by LHC experimentalists in their analyses and

may well serve as a model in other branches of science which have need to compare data and complex simulations

Poster Session 1 - Tableau: 7 / 195

The CMS User Analysis Farm at Fermilab

H. Wenzel¹ ; I. Mandrichenko¹ ; J. Kaiser¹ ; L. Giacchetti¹ ; N. Ratnikova¹

¹ *FERMILAB*

US-CMS is building up expertise at regional centers in preparation for analysis of LHC data. The User Analysis Farm (UAF) is part of the Tier 1 facility at Fermilab. The UAF is being developed to support the efforts of the Fermilab LHC Physics Center (LPC) and to enable efficient analysis of CMS data in the US.

The support, infrastructure, and services to enable a local analysis community at a computing center which is remote from the physical detector and the majority of the collaboration present unique challenges.

The current UAF is a farm running the LINUX operating system providing interactive and batch computing for users. Load balancing, resource and process management are realized with FBSNG, the batch system developed at Fermilab. Over the course of the next three years the UAF must grow in size and functionality, while continuing to support simulated analysis activities and test beam applications.

In this presentation we will describe the development of the current cluster, the technology choices made, the services required to support regional analysis activities, and plans for the future.

Distributed Computing Services / 184

The Clarens Grid-enabled Web Services Framework: Services and Implementation

A. Anjum¹ ; C. Steenberg² ; F. van Lingen³ ; H. Newman⁴ ; I. Legrand⁵ ; J. Bunn³ ; M. Thomas² ; T. Azim¹

¹ *National University of Science and Technology, Rawalpindi*

² *California Institute of Technology*

³ *CALIFORNIA INSTITUTE OF TECHNOLOGY*

⁴ *Caltech*

⁵ *CALTECH*

Auteur(s) contact: conrad@hep.caltech.edu, fvlingen@caltech.edu

Clarens enables distributed, secure and high-performance access to the worldwide data storage, compute, and information Grids being constructed in anticipation of the needs of the Large Hadron Collider at CERN. We report on the rapid progress in the development of a second server implementation in the Java language, the evolution of a peer-to-peer network of Clarens servers, and general improvements in client and server implementations.

Services that are implemented at this time include read/write file access, service lookup and discovery, configuration management, job execution, Virtual Organization Management, an LHCb Information Service, as well as web service interfaces to POOL replica location and metadata catalogs, MonaLISA monitoring information, CMS MCRunjob workflow management, BOSS job monitoring and bookkeeping, Sphinx job scheduler and Chimera virtual data systems.

Commodity web service protocols allows a wide variety of computing platforms and applications to be used to securely access Clarens services, including a standard web browser, Java applets and stand-alone applications, the ROOT data analysis package, as well as libraries that provide programmatic access from the Python, C/C++ and Java languages.

Poster Session 1 - Tableau: 28 / 390

The Condor based CDF CAF

E. Lipeles¹ ; F. Wuerthwein² ; I. Sfiligoi³ ; M. Neubauer⁴

¹ UNIVERSITY OF CALIFORNIA SAN DIEGO

² UCSD

³ INFN Frascati

⁴ University of California, San Diego

Auteur(s) contact: fkw@fnal.gov

The CDF Analysis Facility (CAF) has been in use since April 2002 and has successfully served 100s of users on 1000s of CPUs. The original CAF used FBSNG as a batch manager. In the current trend toward multisite deployment, FBSNG was found to be a limiting factor, so the CAF has been reimplemented to use Condor instead. Condor is a more widely used batch system and is well integrated with the emerging grid tools. One of the most useful being the ability to run seamlessly on top of other batch systems. The transition has brought us a lot of additional benefits, such as ease of installation, fault tolerance and increased manageability of the cluster. The CAF infrastructure has also been simplified a lot since Condor implements a number of features we had to implement ourselves with FBSNG. In addition, our users have found that Condor's fair share mechanism provides a more equitable and predictable distribution of resources. In this talk the Condor based CAF will be presented, with particular emphasis on the changes needed to run with Condor, the problems found during and the advantages gained by the transition. Some background and the plans for the future, as well as results from Condor scalability tests will also be presented.

Poster Session 1 - Tableau: 29 / 130

The Configurations Database Challenge in the ATLAS DAQ System

A. AMORIM¹ ; A. Kazarov² ; D. BURCKHART-CHROMEK³ ; D. KLOSE⁴ ; D. Liko⁵ ; E. BADESCU⁶ ; I. ALEXANDROV⁷ ; I. Soloviev⁸ ; J. FLAMMER³ ; L. MAPELLI³ ; L. PEDRO⁴ ; M. CAPRINI⁶ ; M. DOBSON³ ; M. Mineev⁹ ; N. FIUZA DE BARROS⁴ ; R. JONES³ ; S. KOLOS⁵ ; S. KOROBOV⁷ ; V. Kotov⁹ ; Y. RYABOV¹⁰

¹ *Faculdade de Ciencias, Lisbon, Portugal*

² *Petersburg Nuclear Physics Institute (PNPI)*

³ *CERN, Geneva, Switzerland*

⁴ *Universidade de Lisboa, Portugal*

⁵ *CERN*

⁶ *National Institute of Physics and Nuclear Engineering, Bucharest, Romania*

⁷ *JINR, Dubna, Russia*

⁸ *CERN/PNPI*

⁹ *Joint Institute for Nuclear Research (JINR)*

¹⁰ *PNPI, St-Petersburg, Russia*

Auteur(s) contact: igor.soloviev@cern.ch

The ATLAS data acquisition system uses the database to describe configurations for different types of data taking runs and different sub-detectors. Such configurations are composed of complex data objects with many inter-relations. During the DAQ system initialisation phase the configurations database is simultaneously accessed by a large number of processes. It is also required that such processes be notified about database changes that happen during or between data-taking runs.

The paper describes the architecture of the configurations database. It presents the set of graphical tools which are available for the database schema design and the data editing. The automatic generation of data access libraries for C++ and Java languages is also described. They provide the programming interfaces to access the database either via a common file system or via remote database servers, and the notification mechanism on data changes.

The paper presents results of recent performance and scalability tests, which allow a conclusion to be drawn about the applicability of the current configurations database implementation in the future DAQ system.

Poster Session 2 - Tableau: 71 / 372

The D0 Virtual Center and planning for large scale computing

A. Boehnlein¹ ; D0 Computing Planning Board for the²

¹ *FERMI NATIONAL ACCELERATOR LABORATORY*

² *D0 Collaboration*

Auteur(s) contact: cope@fnal.gov

The D0 experiment relies on large scale computing systems to achieve her physics goals. As the experiment lifetime spans, multiple generations of computing hardware, it is fundamental to make projective models in to use available resources to meet the anticipated needs. In addition, computing resources can be supplied as in-kind contributions by collaborating institutions and countries, however, such resources typically require scheduling, thus adding another dimension for planning. In addition, to avoid over-subscription of the resources, the experiment has to be educated

on the limitations and trade-offs for various computing activities to enable the management to prioritize. We present the metrics and mechanisms used for planning and discuss the uncertainties and unknowns, as well as some of the mechanisms for communicating the resource load to the stakeholders.

In order to correctly account for in-kind contributions of remote computing, D0 uses the concept of a Virtual Center, in which all of the costs are estimated as if the computing were located at solely at FNAL. In contrast to other such models in common use, D0 accounts for contributions based on computer usage rather than strictly on money spend on hardware. This gives incentive to achieve the maximum efficiency of the systems as well as encouraging active participation in the computing model by collaborating institutions. This method of operation leverages a common tool and infrastructure base for all production-type activities.

Online Computing / 119

The DAQ system for the Fluorescence Detectors of the Pierre Auger Observatory

A. KOPMANN¹ ; H.-J. Mathes² ; O. MARTINEAU³ ; S. Argiro⁴

¹ *Forschungszentrum Karlsruhe, Germany*

² *FORSCHUNGSZENTRUM KARLSRUHE, INSTITUT FÜR KERNPHYSIK*

³ *Forschungszentrum Karlsruhe, Germany*

⁴ *INFN Sezione Torino*

Auteur(s) contact: hermann-josef.mathes@cern.ch

S.Argiro⁽¹⁾, A. Kopmann⁽²⁾, O.Martineau⁽²⁾, H.-J. Mathes⁽²⁾
for the Pierre Auger Collaboration

(1) INFN, Sezione Torino

(2) Forschungszentrum Karlsruhe

The Pierre Auger Observatory currently under construction in Argentina will investigate extensive air showers at energies above 10^{18} eV. It consists of a ground array of 1600 Cherenkov water detectors and 24 fluorescence telescopes to discover the nature and origin of cosmic rays at these ultra-high energies.

The ground array is overlooked by 4 different fluorescence buildings which are equipped with 6 telescopes each. An independent local data acquisition (DAQ) is running in each building to readout 480 channels per telescope. In addition, a central DAQ merges data coming from the water detectors and all fluorescence buildings.

The system architecture follows the object oriented paradigm and has been implemented using several of the most widespread open source tools for interprocess communication, data storage and user interfaces.

Each local DAQ is connected with further sub-systems for calibration, for monitoring of atmospheric parameters and slow control. The latter is responsible for general safety functions and the experiment control.

After a prototype phase to validate the system concept the Observatory is taking data in the final setup since September 2003. The data taking will continue during the construction phase and the integration of all sub-systems.

We present the design and the present status of the system currently running in two different buildings with a total of 8 telescopes installed.

Online Computing / 477

The DZERO Run II Level 3 Trigger and Data Acquisition System

A Haas¹ ; D Chapin² ; D Cutts² ; G Brooijmans³ ; G. Watts⁴ ; R Rechenmacher⁵ ; R. HAUSER⁶ ; S Snyder⁷

¹ *University of Washington*

² *Brown University*

³ *Columbia University*

⁴ *UNIVERSITY OF WASHINGTON*

⁵ *Fermi National Accelerator Laboratory*

⁶ *Michigan State University, USA*

⁷ *Brookhaven National Laboratory*

Auteur(s) contact: dpchapin@fnal.gov, gwatts@uw.edu

The DZERO Level 3 Trigger and Data Acquisition (L3DAQ) system has been running continuously since Spring 2002. DZERO is located at one of the two interaction points in the Fermilab Tevatron Collider. The L3DAQ moves front-end readout data from VME crates to a trigger processor farm. It is built upon a Cisco 6509 Ethernet switch, standard PCs, and commodity VME single board computers. We will report on operating experience, performance, and upgrades. In particular, issues related to hardware quality, networking and security, and an expansion of the trigger farm will be discussed.

Poster Session 3 - Tableau: 52 / 283

The Description of the Atlas Detector

A. NAIRZ¹ ; A. VANIACHINE² ; A. ZALITE³ ; B. SELIGMAN⁴ ; C. BOURDARIOS⁵ ; D. COSTANZO⁶ ; D. POMAREDE⁷ ; G. GORFINE⁸ ; G. UNAL⁵ ; I. TRIGGER¹ ; J. BOUDREAU⁹ ; J. TOTH¹⁰ ; K.A. ASSAMAGAN¹¹ ; M. LELCHUK⁴ ; P. Strizeneč¹² ; S. BARANOV¹³ ; S. Spagnolo¹⁴ ; T. KITTELMANN¹⁵ ; V. TSULAIA⁹ ; Y. HASEGAWA¹⁶

¹ *CERN, Geneva, Switzerland*

² *ANL, USA*

³ *Petersburg Nuclear Physics Institute, Russia*

⁴ *Nevis Laboratory, USA*

⁵ *LAL-Orsay, France*

⁶ *LBL, USA*

⁷ *CEA/DAPNIA, Saclay, France*

⁸ *NIKHEF, The Netherlands*

⁹ *Pittsburgh University, USA*

¹⁰ *KFKI-RMKI, Hungary*

¹¹ *BNL, USA*

¹² *IEP SAS Slovakia*

¹³ *Albert-Ludwigs University, Germany*

¹⁴ *DIP. FISICA E INFN, LECCE*

¹⁵ *Niels Bohr Institute, Denmark*

¹⁶ *Shinshu University, Japan*

Auteur(s) contact: stefania.spagnolo@le.infn.it

The ATLAS Detector consists of several major subsystems: an inner detector composed of pixels, microstrip detectors and a transition radiation tracker; electromagnetic and hadronic calorimetry, and a muon spectrometer. Over the last year, these systems have been described in terms of a set of geometrical primitives known as GeoModel. Software components for detector description interpret structured data from a relational database and build from that a complete description of the detector. This description is now used in the Geant-4 based simulation program and also for reconstruction. Detector-specific services that are not handled in a generic way (e.g strip pitches and calorimetric tower boundaries) are added as an additional layer which is synched to the raw geometry. Detector misalignments may also be fed through the model to both simulation and reconstruction. Visualization of the detector geometry is accomplished through Open Inventor and its HEPVis extensions. The ATLAS geometry system in the last year has undergone extensive visual debugging, and experience with the new system has been gained not only through the data challenge but also through the combined test beam. This talk gives an overview of the ATLAS detector description and discusses operational experience with the system in the data challenges and combined test beam.

Poster Session 1 - Tableau: 17 / 304

The Design, Installation and Management of a Tera-Scale High Throughput Cluster for Particle Physics Research

A. Martin¹

¹ QUEEN MARY, UNIVERSITY OF LONDON

Auteur(s) contact: a.j.martin@qmul.ac.uk

We describe our experience in building a cost efficient High Throughput Cluster (HTC) using commodity hardware and free software within a university environment. Our HTC has a modular system architecture and is designed to be upgradable. The current, second phase configuration, consists of 344 processors and 20 Tbyte of RAID storage.

In order to rapidly install and upgrade software, we have developed automatic remote system installation and configuration tools to deploy standard software configurations on individual machines. To efficiently manage machines we have written a custom cluster configuration database. This database is used to track all hardware components in the cluster, the network and power distribution and the software configuration. Access to this database and the cluster performance and monitoring systems is provided by a web portal, which allows efficient remote management in our low-manpower environment.

We describe the performance of our system under a mixed load of scalar and parallel tasks and discuss future possible improvements.

Plenary / 525

The Evolution of Computing: Slowing down? Not Yet!

A. SUTHERLAND^{None}

Dr Sutherland will review the evolution of computing over the past decade, focusing particularly on the development of the database and

middleware from client server to Internet computing.

But what are the next steps from the perspective of a software company? Dr Sutherland will discuss the development of Grid as well as the future applications revolving around collaborative working, which are appearing as the next wave of computing applications.

Plenary / 509

The Evolving Wide Area Network Infrastructure in the LHC era

P. CLARKE^{None}

The global network is more than ever taking its role as the great “enabler” for many branches of science and research. Foremost amongst such science drivers is of course the LHC/LCG programme, although there are several other sectors with growing demands of the network.

Common to all of these is the realisation that a straightforward over provisioned best efforts wide area IP service is probably not enough for the future.

This talk will summarise the needs of several science sectors, and the advances being made to exploit the current best efforts infrastructure. It will then describe current projects aimed at provisioning “better than best efforts” services (such bandwidth on demand), the global optical R&D testbeds and the strategy of the research network providers to move towards hybrid multi-service networks for the next generation of the global wide area production network.

Poster Session 3 - Tableau: 57 / 391

The FEDRA - Framework for Emulsion Data Reconstruction and Analysis in OPERA experiment.

V. Tioukov¹

¹ INFN NAPOLI

Auteur(s) contact: valeri@na.infn.it

OPERA is a massive lead/emulsion target for a long-baseline neutrino oscillation search. More than 90% of the useful experimental data in OPERA will be produced by the scanning of emulsion plates with the automatic microscopes. The main goal of the data processing in OPERA will be the search, analysis and identification of primary and secondary vertexes produced by neutrino in lead-emulsion target.

The volume of middle and high-level data to be analysed and stored is expected to be of the order of several Gb per event. The storage, calibration, reconstruction, analysis and visualization of this data is the task of FEDRA - system written in C++ and based on ROOT framework. The system is now actively used for processing of test beams and simulation data. Several interesting algorithmic solutions permits us to make very effective code for fast pattern

recognition in heavy signal/noise conditions. The system consists of the storage part, intercalibration and segments linking part, track finding and fitting, vertex finding and fitting and kinematical analysis parts. Kalman Filtering technique is used for tracks&vertex fitting. ROOT-based event display is used for interactive analysis of the special events.

Core Software / 400

The FreeHEP Java Library Root IO package

M. Donszelmann¹ ; M. Turri¹ ; T. Johnson¹ ; V. Serbo¹

¹ SLAC

The FreeHEP Java library contains a complete implementation of Root IO for Java. The library uses the “Streamer Info” embedded in files created by Root 3.x to dynamically create high performance Java proxies for Root objects, making it possible to read any Root file, including files with user defined objects. In this presentation we will discuss the status of this code, explain its implementation and demonstrate performance using benchmark comparisons to standard Root IO. We will also describe recently added support for reading files remotely using rootd and xrootd protocols.

We will also show some uses of this library, including using JAS3 to analyze Root data, using the WIRED event display to visualize data from Root files and using rootd and Java servlet technology to make live plots web accessible - with examples from GLAST and BaBar. We will also explain how you can trivially make your own root data web-accessible using the AIDA Tag Library and Jakarta Tomcat.

Distributed Computing Systems and Experiences / 357

The GANGA user interface interface for physics analysis on distributed resources

A. Maier¹ ; A. Soroko² ; C.L. Tan³ ; D. Adams⁴ ; G.N. Patrick⁵ ; J. Martyniak⁶ ; J. Moscicki¹ ; K. Harrison⁷ ; P. Charpentier¹ ; P. Mato¹ ; R. Jones⁸ ; U. Egede⁹

¹ CERN

² Oxford University

³ Birmingham University

⁴ BNL

⁵ RAL

⁶ Imperial College London

⁷ Cambridge University

⁸ Lancaster University

⁹ IMPERIAL COLLEGE LONDON

Auteur(s) contact: u.egede@imperial.ac.uk

Any physicist who will analyse data from the LHC experiments will have to deal with data and computing resources which are distributed across multiple locations and with different access methods. GANGA helps the end user by tying in specifically to the

solutions for a given experiment ranging from specification of data to retrieval and post-processing of produced output. For LHCb and ATLAS the main goal is to assist in running jobs based on the Gaudi/Athena C++ framework. GANGA is written in Python and presents the user with a single GUI rather than a set of different applications. It interacts with external resources like experiments bookkeeping databases, job configuration, and Grid submission systems through plug-able modules. The user is upon start-up presented with a list of templates for common analysis tasks and GANGA persists information about ongoing tasks between invocations. GANGA can also be used through a command line interface that has a tight connection to the GUI to ease the transition from one to the other. Examples will be presented that demonstrates the integration into the distributed analysis systems of the LHCb and ATLAS experiments as used during their 2004 data challenges.

Event Processing / 279

The GeoModel Toolkit for Detector Description

J. Boudreau¹ ; V. Tsulaia¹

¹ UNIVERSITY OF PITTSBURGH

Auteur(s) contact: vakhtang.tsulaia@cern.ch

The GeoModel toolkit is a library of geometrical primitives that can be used to describe detector geometries. The toolkit is designed as a data layer, and especially optimized in order to be able to describe large and complex detector systems with minimum memory consumption. Some of the techniques used to minimize the memory consumption are: shared instancing with reference counting, compressed representations of Euclidean transformations, special nodes which encode the naming of volumes without storing name-strings, and, especially, parameterization through embedded symbolic expressions of transformation fields. A faithful representation of a GeoModel description can be transferred to Geant4, and, we predict, to other engines that simulate the interaction of particles with matter. GeoModel comes with native capabilities for geometry clash detection and for material integration. It's only external dependencies are upon CLHEP.

This talk describes this toolkit for the first time in a public forum.

Poster Session 3 - Tableau: 36 / 354

The Geometry Package for the Pierre Auger Observatory

L. Nellen¹ ; L. Prado Jr² ; S. Argiro³ ; T. Paul⁴ ; T. Porter⁵

¹ I. DE CIENCIAS NUCLEARES, UNAM

² State University of Campinas

³ University of Torino

⁴ Northeastern University

⁵ Louisiana State University

The Pierre Auger Observatory consists of two sites with several semi-autonomous detection systems. Each component, and in some cases each event, provides a preferred coordinate system for simulation and analysis. To avoid a proliferation of coordinate systems in the offline software of the Pierre Auger Observatory, we have developed a

geometry package that allows the treatment of fundamental geometrical objects in a coordinate-independent way. This package makes transformations between coordinate systems transparent to the user, without taking the control about the internal representation completely from the user.

The geometry package allows easy combination of the results from different sub-detectors, at the same time as ensuring that effects like the earth curvature, which is non-negligible on the scale of a single Auger site, are dealt with properly.

The internal representations used are Cartesian. For interfacing, including I/O, the package includes support for Cartesian coordinates, geodetic (latitude/longitude and UTM), and astrophysical coordinate systems.

Grid Security / 249

The GridSite authorization system

A. McNab¹ ; S. Kaushal¹

¹ UNIVERSITY OF MANCHESTER

Auteur(s) contact: andrew.mcnab@man.ac.uk

We describe the GridSite authorization system, developed by GridPP and the EU DataGrid project for access control in High Energy Physics grid environments with distributed virtual organizations. This system provides a general toolkit of common functions, including the evaluation of access policies (in GACL or XACML), the manipulation of digital credentials (X.509, GSI Proxies or VOMS attribute certificates) and utility functions for protocols such as HTTP.

GridSite also provides a set of extensions to the Apache web server to permit it to function in a Grid security environment, including access control, fileserver / webservice management and a lightweight Virtual Organization service.

Using Apache as an example, we explain how Grid security can be added to an existing service using our toolkit. We then outline some of the other uses to which components have been put in the deployed Grids of GridPP, the EU DataGrid and the LHC Computing Grid.

Online Computing / 411

The High Level Filter of the H1 Experiment at HERA

A. Campbell¹ ; M. Vorobiev² ; S. Levonian¹

¹ DESY

² DESY,ITEP

We present the scheme in use for online high level filtering, event reconstruction and classification in the H1 experiment at HERA since 2001.

The Data Flow framework (presented at CHEP2001) will be reviewed. This is based on CORBA for all data transfer,

multi-threaded C++ code to handle the data flow and synchronisation and fortran code for reconstruction and event selection. A controller written in python provides setup,initialisation and process management. Specialised java programs provide run control and online access to and display of histograms. A C++ logger program provides central logging of standard printout from all processes.

We show how the system handles online preparation and update of detector calibration and beam parameter data. Newer features are the selection of rare events for the online event display and the extension to multiple input sources and output channels.

We discuss how the system design provides automatic recovery from various failures and show the overall and long term performance.

In addition we present the framework of event selection and classification and the features it provides.

Event Processing / 444

The High Level Trigger software for the CMS experiment

C. Delaere¹ ; O. van der Aa¹

¹ *INSTITUT DE PHYSIQUE NUCLEAIRE, UNIVERSITE CATHOLIQUE DE LOUVAIN*

Auteur(s) contact: olivier.van.der.aa@cern.ch

The observation of Higgs bosons predicted in supersymmetric theories will be a challenging task for the CMS experiment at the LHC, in particular for its High Level trigger (HLT). A prototype of the High Level Trigger software to be used in the filter farm of the CMS experiment and for the filtering of monte carlo samples will be presented. The implemented prototype heavily uses recursive processing of a HLT tree and allows dynamic trigger definition. Firstly the general architecture and design choices as well as the timing performance of the system will be reviewed in the light of the DAQ constrains. Secondly, specific trigger implementations in the context of the object-oriented Reconstruction for CMS Analysis (ORCA) software will be detailed. Finally, the analysis for the selection of a CP even Higgs decaying in tau pairs will be presented. The Aforementioned analysis will illustrate the importance of the trigger strategies required to achieve the various physics analysis in CMS.

Plenary / 528

The IBM Research Global Technology Outlook

D. MCQUEENEY^{None}

The Global Technology Outlook (GTO) is IBM Research's projection of the future for information technology (IT). The GTO identifies progress and trends in key indicators such as raw computing speed, bandwidth, storage, software technology, and business modeling. These new technologies have the

potential to radically transform the performance and utility of tomorrow's information processing systems and devices, ultimately creating new levels of business value.

Plenary / 522

The LCG Project - Preparing for Startup

L. ROBERTSON^{None}

The talk will cover briefly the current status of the LHC Computing Grid project and will discuss the main challenges facing us as we prepare for the startup of LHC.

Poster Session 3 - Tableau: 41 / 384

The LCG Savannah software development portal

Auteur(s): Y. Perrin¹

Co-auteur(s) D. Feichtinger²; F. Orellana³; M. Roy⁴

¹ CERN

² PSI

³ Univ. of Geneva

⁴ FSF France

A web portal has been developed, in the context of the LCG/SPI project, in order to coordinate workflow and manage information in large software projects. It is a development of the GNU Savannah package and offers a range of services to every hosted project: Bug / support / patch trackers, a simple task planning system, news threads, and a download area for software releases. Features and functionality can be fine-tuned on a per project basis and the system displays content and grants permissions according to the user's status (project member, other Savannah user, or visitor). A highly configurable notification system is able to channel tracker submissions to developers in charge of specific project modules.

The portal is based on the GNU Savannah package which is now developed as 'Savane' by the Free Software Foundation of France. It is a descendant of the well known SourceForge-2.0 software. The original trackers were contributed to the open source community by XEROX, which uses a similar system for their internal software development. Several features and extensions were introduced in a collaboration of LCG/SPI with the current main developer of Savannah to adapt the software for use at CERN and the results were given back to the open source. Cern Savannah currently provides services to more than 600 users in 90 projects.

Poster Session 2 - Tableau: 8 / 436

The LCG-AliEn interface, a realization of a MetaGrid system

A. Peters¹; P. Buncic¹; P. Cerello²; P. Saiz¹; S. Bagnasco²

¹ *CERN*

² *INFN Torino*

AliEn (ALICE Environment) is a GRID middleware developed and used in the context of ALICE, the CERN LHC heavy-ion experiment. In order to run Data Challenges exploiting both AliEn “native” resources and any infrastructure based on EDG-derived middleware (such as the LCG and the Italian GRID.IT), an interface system was designed and implemented; some details of a prototype were already presented at CHEP2003. In the spring of 2004 an ALICE Data Challenge began with the simulated data production on this multiple infrastructure, thus qualifying as the first large production carried out transparently making use of very different middleware system. This system is a practical realisation of the “federated” or “meta-” grid concept, and it has been successfully tested in a very large production. This talk reports about new developments of the interface system, the successful DC running experience, the advantages and limitations of this concept, the plans for the future and some lessons learned.

Online Computing / 74

The LHCb Configuration Database

B. Jost¹ ; C. Gaspar¹ ; E. Van Herwijnen¹ ; L. Abadie¹ ; N. Neufeld¹ ; R. Jacobsson¹

¹ *CERN*

The aim of the LHCb configuration database is to store all the controllable devices of the detector. The experiment’s control system (that uses PVSS) will configure, start up and monitor the detector from the information in the configuration database. The database will contain devices with their properties, connectivity and hierarchy. The ability to rapidly store and retrieve huge amounts of data, and the navigability between devices are important requirements. We have collected use cases to ensure the completeness of the design. Using the entity relationship modeling technique we describe the use cases as classes with attributes and links. We designed the schema of the tables using the relational diagrams. This methodology has been applied to describe and store the connectivity of the devices in the TFC (switches) and DAQ system. Other parts of the detector will follow later.

The database has been implemented using Oracle to benefit from central CERN database support. The project also foresees the creation of tools to populate, maintain, and configure the configuration database. To communicate between the control system and the database we have developed a system which sends queries to the database and displays the results in PVSS. This database will be used in conjunction with the configuration database developed by the CERN JCOP project for PVSS.

Wide Area Networking / 96

The Network Security Protection System at IHEP-Net

Auteur(s): L. Ma¹

Co-auteur(s) B. Liu² ; C. Wu² ; C. Yu² ; D. An² ; G. Chen² ; R. Liu² ; R. Xu²

¹ *INSTITUTE OF HIGH ENERGY PHYSICS*

² *Institute of High Energy Physics*

Auteur(s) contact: ma@ihep.ac.cn

Network security at IHEP is becoming one of the most important issues of computing environment. To protect its computing and network resources against attacks and viruses from outside of the institute, security measures to combat these are implemented. To enforce security policy the network infrastructure was re-configured to one intranet and two DMZ areas. New rules to control the access between intranet and DMZ areas are applied. All hosts at IHEP are divided into three types according to their security levels. Hosts of the first type are isolated in the institute and can just access the hosts inside of IHEP. The second type hosts access Internet through NAT. The third type hosts will directly connect to outside. An intrusion detection system works with firewall so that all packets from outside IHEP are checked and filtered. Access from outside will go through firewall or VPN. In order to prevent virus spread at IHEP and reduce the number of spam mail we installed a virus filter and spam filter system. All of these measures make the network at IHEP more secure. Attacks, virus and spam mails decrease dramatically.

Distributed Computing Services / 328

The Next Generation Root File Server

A. Dorigo¹ ; A. Hanushevsky² ; F. Furano¹

¹ *INFN, Padova*

² *SLAC*

As the BaBar experiment shifted its computing model to a ROOT-based framework, we undertook the development of a high-performance file server as the basis for a fault-tolerant storage environment whose ultimate goal was to minimize job failures due to server failures. Capitalizing on our five years of experience with extending Objectivity's Advanced Multithreaded Server (AMS), elements were added to remove as many obstacles to server performance and fault-tolerance as possible. The final outcome was xrootd, upwardly and downwardly compatible with the current file server, rootd. This paper describes the essential protocol elements that make high performance and fault-tolerance possible; including asynchronous parallel requests, stream multiplexing, data pre-fetch, automatic data segmenting, and the framework for a structured peer-to-peer storage model that allows massive server scaling and client recovery from multiple failures. The internal architecture of the server is also described to explain how high performance was maintained and full compatibility was achieved. Now in production at Stanford Linear Accelerator Center, Rutherford Appleton Laboratory (RAL), INFN, and IN2P3; xrootd has shown that our design provides what we set out to achieve. The xrootd server is now part of the standard ROOT distribution so that other experiments can benefit from this data serving model within a standard HEP event analysis framework.

Core Software / 353**The Offline Framework of the Pierre Auger Observatory**L. Nellen¹ ; L. Prado Jr² ; S. Argiro³ ; T. Paul⁴ ; T. Porter⁵¹ *I. DE CIENCIAS NUCLEARES, UNAM*² *State University of Campinas*³ *University of Torino*⁴ *Northeastern University*⁵ *Louisiana State University*

The Pierre Auger Observatory is designed to unveil the nature and the origin of the highest energy cosmic rays. Two sites, one currently under construction in Argentina, and another pending in the Northern hemisphere, will observe extensive air showers using a hybrid detector comprising a ground array of 1600 water Cerenkov tanks overlooked by four atmospheric fluorescence detectors. Though the computing demands of the experiment are less severe than those of traditional high energy physics experiments in terms of data volume and detector complexity, the large geographically dispersed collaboration and the heterogeneous set of simulation and reconstruction requirements confronts the offline software with some special challenges.

We have designed and implemented a framework to allow collaborators to contribute algorithms and sequencing instructions to build up the variety of applications they require. The framework includes machinery to manage these user codes, to organize the abundance of user-contributed configuration files, to facilitate multi-format file handling, and to provide access to event and time-dependent detector information which can reside in various data sources. A number of utilities are also provided, including a novel geometry package which allows manipulation of abstract geometrical objects independent of coordinate system choice. The framework is implemented in C++, follows an object oriented paradigm, and takes advantage of some of the more widespread tools that the open source community offers, while keeping the user-side simple enough for C++ non-experts to learn in a reasonable time. The distribution system includes unit and acceptance testing in order to support rapid development of both the core framework and contributed user code. Great attention has been paid to the ease of installation.

Distributed Computing Systems and Experiences / 192**The Open Science Grid (OSG)**R. Pordes¹¹ *FERMILAB***Auteur(s) contact:** ruth@fnal.gov

The U.S.LHC Tier-1 and Tier-2 laboratories and universities are developing production Grids to support LHC applications running across a worldwide Grid computing system. Together with partners in computer science, physics grid projects and running experiments, we will build a common national production grid infrastructure which is open in its architecture, implementation and use.

The OSG model builds upon the successful approach of last year's joint Grid2003 project. The Grid3 shared infrastructure has for over eight months given significant computational resources and throughput to more than six applications, including ATLAS and CMS data challenges, SDSS, LIGO and Biology analyses and computer science demonstrators.

To move towards LHC-scale data management, access and analysis capabilities, we will need to increase the scale, services, and sustainability of the current infrastructure by an order of magnitude. This requires a significant upgrade in its functionalities and technologies.

The OSG roadmap is a strategy and work plan to build the U.S.LHC computing enterprise as a fully usable, sustainable and robust grid, which is part of the LHC global computing infrastructure and open to partners. The approach is to federate with other application communities in the U.S. to build a shared infrastructure open to other sciences and capable of being modified and improved to respond to needs of other applications, including CDF, D0, BaBar and RHIC experiments.

We describe the application driven engineered services of the OSG, short term plans and status, and the roadmap for a consortium, its partnerships and national focus.

Online Computing / 329

The PHENIX Event Builder

D. Winter¹

¹ COLUMBIA UNIVERSITY

Auteur(s) contact: winter@nevis.columbia.edu

The PHENIX detector consists of 14 detector subsystems. It is designed such that individual subsystems can be read out independently in parallel as well as a single unit. The DAQ used to read the detector is a highly-pipelined parallel system. Because PHENIX is interested in rare physics events, the DAQ is required to have a fast trigger, deep buffering, and very high bandwidth.

The PHENIX Event Builder is a critical part of the back-end of the PHENIX DAQ. It is responsible for assembling event fragments from each subsystem into complete events ready for archiving. It allows subsystems to be read out either in parallel or simultaneously and supports a high rate of archiving. In addition, it implements an environment where Level-2 trigger algorithms may be optionally executed, providing the ability to tag and/or filter rare physics events.

The Event Builder is a set of three Windows NT/2000 multithreaded executables that run on a farm of over 100 dual-cpu 1U servers. All control and data messaging is transported over a Foundry Layer2/3 Gigabit switch. Capable of recording a wide range of event sizes from central Au-Au to p-p interactions, data archiving rates of over 400 MB/s at 2 KHz event rates have been achieved in the recent Run 4 at RHIC. Further improvements in performance are expected from migrating to Linux for Run 5.

The PHENIX Event Builder design and implementation, as well as performance and plans for future development, will be discussed.

Poster Session 1 - Tableau: 24 / 495

The Project CampusGrid

O. Schneider¹

¹ *FZK*

Auteur(s) contact: olaf.schneider@iwr.fzk.de

A central idea of Grid Computing is the virtualization of heterogeneous resources. To meet this challenge the Institute for Scientific Computing, IWR, has started the project CampusGrid. Its medium term goal is to provide a seamless IT environment supporting the on-site research activities in physics, bioinformatics, nanotechnology and meteorology. The environment will include all kinds of HPC resources: vector computers, shared memory SMP servers and clusters of commodity components as well as a shared high-performance storage solution. After introducing the general ideas the talk will inform about the current project status and scheduled development tasks. This is associated with reports on other activities in the fields of Grid computing and high performance computing at IWR.

Poster Session 3 - Tableau: 10 / 303

The ROOT Linear Algebra

E. Offermann¹ ; F. Rademakers¹ ; R. Brun¹

¹ *CERN*

Auteur(s) contact: rene.brun@cern.ch

The ROOT linear algebra package has been invigorated . The hierarchical structure has been improved allowing different flavors of matrices, like dense and symmetric . A fairly complete set of matrix decompositions has been added to support matrix inversions and solving linear equations.

The package has been extensively compared to other algorithms for its accuracy and performance.

In this poster we will describe the structure of the package and several benchmarks obtained with typical linear algebra applications.

Computer Fabrics / 462

The SAMGrid Database Server Component: Its Upgraded Infrastructure and Future Development Path

A. Baranovski¹ ; A. Kreymer² ; A. Lyon² ; A. Sill³ ; F. Ratnikov⁴ ; G. Garzoglio² ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter¹ ; L. Lueking² ; M. Burgon-Lyon⁵ ; M. Leslie⁶ ; R. Herber² ; R. Illingworth² ; R. Kennedy² ; R. St.Denis⁵ ; S. Belforte⁷ ; S. Stonjek⁶ ; S. Veseli¹ ; S. White¹ ; U. Kerzel⁸ ; V. Bartsch⁶ ; W. Merritt²

¹ *Fermilab*

² *FERMI NATIONAL ACCELERATOR LABORATORY*

³ *Texas Tech. University*

⁴ *Rutgers University*

⁵ *Glasgow University*

⁶ *Oxford University*

⁷ *INFN Trieste*

⁸ *Karlsruhe University*

Auteur(s) contact: veseli@fnal.gov

The SAMGrid Database Server encapsulates several important services, such as accessing file metadata and replica catalog, keeping track of the processing information, as well as providing the runtime support for SAMGrid station services. Recent deployment of the SAMGrid system for CDF has resulted in unification of the database schema used by CDF and D0, and the complexity of changes required for the unified metadata catalog has warranted a complete redesign of the DB Server.

We describe here the architecture and features of the new server. In particular, we discuss the new CORBA infrastructure that utilizes python wrapper classes around IDL structs and exceptions. Such infrastructure allows us to use the same code on both server and client sides, which in turn results in significantly improved code maintainability and easier development.

We also discuss future integration of the new server with an SBIR II project which is directed toward allowing the dserver to access distributed databases, implemented in different DB systems and possibly using different schema.

Core Software / 222

The SEAL C++ Reflection System

P. MATO¹ ; S. Roiser¹

¹ *CERN*

Auteur(s) contact: stefan.roiser@cern.ch

The C++ programming language has very limited capabilities for reflection information about its objects. In this paper a new reflection system will be presented, which allows complete introspection of C++ objects and has been developed in the context of the CERN/LCG/SEAL project in collaboration with the ROOT project.

The reflection system consists of two different parts. The first part is a code generator that produces automatically reflection information from existing C++ classes. This generation of the reflection information is done in a non intrusive way, which means that the original C++ classes definition do not need to be changed or instrumented. The second part of the reflection system is able to load/build this information in memory and provides an API to the user.

The user can query reflection information from any C++ class and also

interact generically with the objects, like invocation of functions, setting and getting data members or constructing and deleting objects. When designing the different packages it was taken care of having minimal dependencies on external software and a possibility to port the software to different platforms/compilers.

A quick overview of the current implementation in use by the LCG SEAL and POOL projects will be given. A more detailed description of the new model, which aims to reflect the complete C++ language and to be a common reflection system used also by the ROOT framework, will be given.

Core Software / 175

The SEAL Component Model

L. Tuura¹ ; P. Mato² ; R. Chytracsek²

¹ *Northeastern University*

² *CERN*

Auteur(s) contact: radovan.chytracsek@cern.ch, pere.mato@cern.ch

This paper describes the component model that has been developed in the context of the LCG/SEAL project. This component model is an attempt to handle the increasing complexity in the current data processing applications of LHC experiments. In addition, it should facilitate the software re-use by the integration of software components from LCG and non-LCG into the experiment's applications. The component model provides the basic mechanisms and base classes that facilitate the decomposition of the whole C++ object-oriented application into a number of run-time pluggable software modules with well defined generic behavior, inter-component interaction protocols, run-time configuration and user customization.

This new development is based on the ideas and practical experiences of the various software frameworks in use by the different LHC experiments for several years. The design and implementation choices will be described and the practical experiences and difficulties in adopting this model to existing experiment software systems will be outlined.

Distributed Computing Services / 318

The STAR Unifid Meta-Scheduler project, a front end around evolving technologies for user analysis and data production.

D. Olson¹ ; E. Efstathiadis² ; G. CARCASSI² ; I. Sakrejda³ ; J. Lauret⁴ ; L. Didenko² ; V. Fine²

¹ *Lawrence Berkeley Laboratory*

² *BROOKHAVEN NATIONAL LABORATORY*

³ *NERSC*

⁴ *Brookhaven National Laboratory*

While many success stories can be told as a product of the Grid middleware developments, most of the existing systems relying on workflow and job execution are based on integration of self-contained production systems interfacing with a given scheduling component or portal, or directly uses the base component of the Grid middleware (globus-job-run, globus-job-submit). However, such systems usually do not take advantage of the presence of Resource Manager System (RMS); they hardly allow

for a mix of local RMS and are either Grid or non-grid enabled. We intend to present an approach taking advantage of both worlds.

The STAR Unified Meta-Scheduler (SUMS) project provides users a way to submit jobs on a farm, at a site (multiple pools or farms) or on the Grid without the need to know or adapt to the diversity of technologies and knowledge involved while using multiple LRMS and their specificities. The strategy was adopted in 2002 to shield the users against changes in technologies inherent to the emerging Grid infrastructure and developments.

Java based and taking as input a simple user job description language (U-JDL), SUMS allows connection with multiple (overlapping or not) LRMS and Grid job submission (Condor-G, grid-job-submit, ...) without the need for changing the U-JDL. Fully integrated with the STAR File and Replica Catalog, information providers (load and queue information), SUMS provides a single point of reference for users to migrate from a traditional to a distributed computing environment. Results and the evolutionist architecture of the SUMS will be presented and its future, improvements and evolution will be discussed.

Poster Session 3 - Tableau: 2 / 133

The TAG Collector - A Tool for Atlas Code Release Management

F. Lambert¹ ; J. Collot¹ ; J. Fulachier¹ ; S. Albrand¹

¹ LPSC

Auteur(s) contact: solveig.albrand@lpsc.in2p3.fr

The Tag Collector is a web interfaced database application for release management. The tool is tightly coupled to CVS, and also to CMT, the configuration management tool. Developers can interactively select the CVS tags to be included in a build, and the complete build commands are produced automatically. Other features are provided such as verification of package CMT requirements files, and direct links to the package documentation, making it a useful tool for all ATLAS users.

The software for the Atlas experiment contains about 1 MSLOC. It is organized in over 50 container packages containing about 500 source code packages. One or several developers maintain each package. ATLAS developers are widely distributed geographically.

The Tag Collector was designed and implemented during the summer of 2001, in response to a near crisis situation. It has been in use since September 2001. Until this time the ATLAS librarian constructed a build of the software release after a cascade of e-mails from developers; communicating the correct CVS code repository version tag of their respective packages. This was subject to all sorts of human errors, and inefficient in our multi-time zone environment. In addition, it was difficult to manage the contents of a release. It was all too easy for a prolific developer to introduce a well-intentioned change in his package just before a build, often with unsuspected border effects. Developers were also asking for regular, and frequent developer builds.

The tool has proved extremely successful, and features that are outside the scope of the original design have been requested. Requirements for a new version were collected during 2003, culminating in a formal review in December 2003. The new version is currently being designed. It will be more flexible and easier to maintain.

Poster Session 3 - Tableau: 49 / 109

The Track Extrapolation package in the new ATLAS Tracking Realm

A. Salzburger¹

¹ *UNIVERSITY OF INNSBRUCK*

Auteur(s) contact: andreas.salzburger@cern.ch

The ATLAS reconstruction software requires extrapolation to arbitrary oriented surfaces of different types inside a non-uniform magnetic field. In addition multiple scattering and energy loss effects along the propagated trajectories have to be taken into account. A good performance in respect of computing time consumption is crucial due to hit and track multiplicity in high luminosity events at the LHC and the small time window of the ATLAS high level trigger.

Therefore stable and fast algorithms for the propagation of the track parameters and their associated covariance matrices in specific representations to different surfaces in the detector are required.

The recently developed track extrapolation package inside the new ATLAS offline tracking software is presented. Timing performance studies, integration tests with client algorithms and results on ATLAS 2004 Combined Test Beam data are given.

Event Processing / 387

The Virtual Geometry Model

I. Hrivnacova¹

¹ *IPN, ORSAY, FRANCE*

Auteur(s) contact: ivana.hrivnacova@cern.ch

In order for physicist to easily benefit from the different existing geometry tools used within the community, the Virtual Geometry Model (VGM) has been designed. In the VGM we introduce the abstract interfaces to geometry objects and an abstract factory for geometry construction, import and export. The interfaces to geometry objects were defined to be suitable to describe "geant-like" geometries with a hierarchical volume structure.

The implementation of the VGM for a concrete geometry model represents a small layer between the VGM and the particular native geometry.

At the present time this implementation is provided for the Geant4 and the Root TGeo geometry models.

Using the VGM factory, geometry can first be defined independently from a concrete geometry model, and then built by choosing a concrete instantiation of it. Alternatively, the import function of the VGM factory makes it possible to use VGM directly with native geometries (Geant4, TGeo). The export functions provide conversion into other native geometries or the XML format.

In this way, the VGM surpasses one-directional geometry converters within Geant4 VMC (Virtual Monte Carlo): roottog4 and g4toxml, and automatically provides missing directions: g4toroot, roottoxml. To port a third geometry model, then providing the VGM layer for it is sufficient to obtain all the converters between this third geometry and already ported geometries (Geant4, Root).

The design and implementation of the VGM classes, the status of existing implementations for Geant4 and TGeo, and simple examples of usage will be presented.

Event Processing / 433

The Virtual MonteCarlo : status and applications

A. Fasso¹ ; A. Gheata² ; A. Morsch² ; E. Futo² ; F. Carminati² ; I. Hrivnacova³ ; M. Gheata⁴ ; P. Hristov² ; R. Brun²

¹ SLAC

² CERN

³ IPN Orsay

⁴ Institute of Space Science - Bucharest

Auteur(s) contact: federico.carminati@cern.ch, andreas.morsch@cern.ch

The current major detector simulation programs, i.e. GEANT3, GEANT4 and FLUKA have largely incompatible environments. This forces the physicists willing to make comparisons between the different transport Monte Carlos to develop entirely different programs. Moreover, migration from one program to the other is usually very expensive, in manpower and time, for an experiment offline environment, as it implies substantial changes in the simulation code. To solve this problem, the ALICE Offline project has developed a virtual interface to these three programs allowing their seamless use without any change in the framework, the geometry description or the scoring code. Moreover a new geometrical modeller has been developed in collaboration with the ROOT team, and successfully interfaced to the three programs. This allows the use of one description of the geometry, which can be used also during reconstruction and visualisation. The talk will describe the present status and future plans for the Virtual Monte Carlo. It will also present the capabilities and performance of the geometrical modeller.

Grid Security / 191

The Virtual Organization Membership Service eXtension project (VOX)

A. Heavey¹ ; D. Skow² ; E. Berman¹ ; G. Carcassi³ ; G. Graham² ; I. Fisk² ; J. Kaiser² ; J. Weigand¹ ; L. Bauerdick² ; L. Grundhoefer⁴ ; R. Baker³ ; R. Gardner⁵ ; R. Pordes² ; T. Levshina¹ ; V. Sekhri¹ ; Y. Wu¹

¹ FERMILAB

² Fermilab

³ BNL

⁴ University of Indiana

⁵ University of Chicago

Auteur(s) contact: ian.fisk@cern.ch

Current grid development projects are being designed such that they require end users to be authenticated under the auspices of a “recognized” organization, called a Virtual Organization (VO). A VO must establish resource-usage agreements with grid resource providers. The VO is responsible for authorizing its members for grid computing privileges. The individual sites and resources typically enforce additional layers of authorization.

The VOX project developed at Fermilab is an extension of VOMS, developed jointly for DataTAG by INFN and for DataGrid by CERN. The Virtual Organization Membership Registration Service (VOMRS) is a major component of the VOX project. VOMRS is a service that provides

the means for registering members of a VO, and coordination of this process among the various VO and grid administrators. It consists of a database to maintain user registration and institutional information, a server to handle members' notification and synchronization with various interfaces, web services and a web user interface for the input of data into the database and manipulation of that data.

The VOX project also includes a component for the Site Authorization (SAZ), which allows security authorities at a site to control access to site resources and a component for the Local Resource Administration (LRAS), which associates the VO member with the local account and local resources on a grid cluster.

The current state of deployment and future steps to improve the prototype and implement some new features will be presented.

Poster Session 3 - Tableau: 26 / 263

The WEB interface for the ATLAS/LCG MySQL Conditions Databases and performance constraints in the visualisation of extensive scientific/technical data

A. Amorim¹ ; A. Vaniachine² ; D. Klose³ ; J. Castilho¹ ; L. Pedro¹ ; N. Barros¹ ; T. Franco¹

¹ FACULTY OF SCIENCES OF THE UNIVERSITY OF LISBON

² Argonne National Laboratory

³ Universidade de Lisboa, Portugal

Auteur(s) contact: luis.pedro@cern.ch

A common LCG architecture for the Conditions Database for the time evolving data enables the possibility to separate the interval-of-validity (IOV) information from the conditions data payload. The two approaches can be beneficial in different cases and separation presents challenges for efficient knowledge discovery, navigation and data visualization. In our paper we describe the conditions data browser - CondDBrowser - a tool deployed in ATLAS for scientific analysis and visualization of this data.

A wide availability and access to the overall distributed conditions data repository was achieved through a seamless integration of the IOV and the payload data to the user a unifying web interface that hides the persistency storage details. Another user-friendly feature of the tool is a simplified querying language similar to QBE (Query by Example).

Our case study is based on the web interface developed for the ATLAS/LCG ConditionsDB. The interaction with other payload storage technologies, external to the ConditionsDB, will also be presented. In particular, the integration of the NOVA database technologies.

We will discuss how the information is gathered from the ConditionsDB and the corresponding extensions needed to enable data browsing in the external repositories, how it is organized, and what kind of operations (search and visualization) are allowed. We'll also present how this interface uses the C++ API extending it to a similar PHP interface, that can be used to browse data collected using any of the ConditionsDB implementations.

Performance constraints are also presented and will be discussed in detail.

Poster Session 3 - Tableau: 55 / 423

The ZEUS Global Tracking Trigger

M. Sutton¹

¹ UNIVERSITY COLLEGE LONDON

Auteur(s) contact: sutt@cern.ch

The design, implementation and performance of the ZEUS Global Tracking Trigger (GTT) Forward Algorithm is described. The ZEUS GTT Forward Algorithm integrates track information from the ZEUS Micro Vertex Detector (MVD) and forward Straw Tube Tracker (STT) to provide a picture of the event topology in the forward direction ($1.5 < \eta < 3$) of the ZEUS detector. This region is particularly challenging because of inhomogeneities in the solenoid magnetic field, and the high occupancies in the forward direction from beam-gas interactions and secondary scatters with the ZEUS beampipe. The forward algorithm is distinct from the GTT barrel algorithm, but will run in parallel on the GTT CPU farm. To avoid unacceptable deadtime in the ZEUS readout system, the forward algorithm processing must be compliant with the strict requirements of the ZEUS trigger system. The current status of the integration with the ZEUS DAQ and trigger systems is also reviewed.

Event Processing / 421

The ZEUS Global Tracking Trigger Barrel Algorithm

M. Sutton¹

¹ UNIVERSITY COLLEGE LONDON

Auteur(s) contact: sutt@cern.ch

The current design, implementation and performance of the ZEUS global tracking trigger barrel algorithm are described. The ZEUS global tracking trigger integrates track information from the ZEUS central tracking chamber (CTD) and micro vertex detector (MVD) to obtain a global picture of the track topology in the ZEUS detector at the second level trigger stage. Algorithm processing is performed on a farm of Linux PCs and, to avoid unacceptable deadtime in the ZEUS readout system, must be completed within the strict requirements of the ZEUS trigger system. The GTT plays a vital role in the selection of good physics events and the rejection of non-physics background within the very harsh trigger environment provided by the upgraded HERA collider. The GTT barrel algorithm greatly improves the vertex resolution and the track finding efficiency of the ZEUS second level trigger while the mean event processing latency and throughput are well within the trigger requirements. Recent running experience with HERA production luminosity is briefly discussed.

Core Software / 356

The ZOOM Minimization Package

D. Sachs¹ ; M. Fischler²

¹ *Fermilab*

² *FERMILAB*

Auteur(s) contact: mf@fnal.gov

A new object-oriented Minimization package is available via the ZOOM cvs repository. This package, designed for use in HEP applications, has all the capabilities of Minuit, but is a re-write from scratch, adhering to modern C++ design principles.

A primary goal of this package is extensibility in several directions, so that its capabilities can be kept fresh with as little maintenance effort as possible. These flexibility goals have been met, as demonstrated by extensions of the package to add new types of termination conditions, new domains and restrictions on the solution space, and a new minimization algorithm. Each of these extensions was straightforward to implement.

The object-oriented design style also has several advantages at the user level. One such advantage is the ability to consider several problems simultaneously within a single program. Another is that we can make the Minimizer objects easy to coordinate the Minimizer with the use of other products. To verify that this goal is met, we demonstrate examples of using the Minimizer in the context of a Root application, and in an application using the “R” statistical analysis environment and language.

We compare and contrast this package with other free C++ Minimization packages suitable for HEP (most of which have origins in Minuit). Following Minuit overly precisely makes it difficult to design an object oriented package without undue distortions. This package is distinguished by the priority that was assigned to C++ design issues, and the focus on producing an extensible system that will resist becoming obsolete.

Poster Session 3 - Tableau: 67 / 73

The application of PowerPC/VxWorks to the read-out subsystem of the BESIII DAQ

N. tao¹ ; YP. chu¹ ; Z. JINGWEI²

¹ *IHEP*

² *IHEP, Beijing, China*

Auteur(s) contact: ninatao@mail.ihep.ac.cn

This article describes the simulation of the read-out subsystem which will be subject to the BESIII data acquisition system. According to the purpose of the BESIII, the event rate will be about 4000Hz, and the data rate up to 50Mbytes/sec after Level 1 trigger. The read-out subsystem consists of some read-out crates and read-out computer whose principle function is to collect event data from the front-end electronics after Level 1 trigger, to transfer data fragments from each VME read-out crate to online computer farm through two levers of computer pre-processing and high speed network transmission. The read-out implementation is based on commercial single board computer MVME5100 running VxWorks operating system.

The article outlines the structure of the simulative platform which included hardware components and software components. It puts emphasis on the framework of the read-out subsystem, data process flow and test method.

Especially, it enumerates key technologies in the process of design and analyses of the test result. In addition, results which summaries the performance of the single board computer from the data transferring aspects will be presented.

Key word: BESIII read-out subsystem MVME5100 VxWorks VMEbus DMA read-out computer

Distributed Computing Systems and Experiences / 440

The architecture of the AliEn system

A. Peters¹ ; J.F. Grosse-Oetringhaus¹ ; P. Buncic¹ ; P. Saiz¹

¹ CERN

Auteur(s) contact: predrag.buncic@cern.ch

AliEn (ALICE Environment) is a Grid framework developed by the Alice Collaboration and used in production for almost 3 years. From the beginning, the system was constructed using Web Services and standard network protocols and Open Source components. The main thrust of the development was on the design and implementation of an open and modular architecture. A large part of the component came from state-of-the-art modules available in the Open Source domain. Thus, in a very short time, the ALICE experiment had a prototype Grid that, while constantly evolving, has allowed large distributed simulation and reconstruction vital to the design of the experiment hardware and software to be performed with very limited manpower. This proved to be the correct path to which many Grid project and initiatives are now converging. The architecture of AliEn inspired the ARDA report and subsequently AliEn provided the foundation of components for the first EGEE prototype. This talk presents the architecture of the original AliEn system, describes its evolution. A critical review of the major technology choices, their implementation and the development process is also presented.

Poster Session 2 - Tableau: 7 / 167

The deployment mechanisms for the ATLAS software.

A. De Salvo¹ ; C. ARNAULT² ; G. Rybkine³ ; S. George⁴

¹ INFN

² CNRS

³ INR

⁴ RHUL

Auteur(s) contact: arnault@lal.in2p3.fr

One of the most important problems in software management of a very large and complex project such as Atlas is how to deploy the software on the running sites. By running sites we include computer sites ranging from computing centers in the usual sense down to individual laptops but also the computer elements of a computing grid organization. The deployment activity consists in constructing a well defined representation of the states of the working software (known as releases), and transporting them to the target sites, in such a way that the installation process can be entirely automated and can take care of discovering the context and adapting itself to it. A set of tools based on both CMT - the basic configuration management tool of ATLAS - and Pacman has been developed. The resulting mechanism now supports the systematic production of distribution kits for various binary conditions of every release, the partial or complete automatic installation of kits on any site and the running of test suites to validate the installed kits. This mechanism is meant to be fully compliant with the Grid requirements and has been tested in the context of LCG. Several issues related with the constraints on the target system, or with the incremental updates of the installation still need to be studied and will be discussed.

Distributed Computing Systems and Experiences / 210

The evolution of the distributed Event Reconstruction Control System in BaBar

A. Ceseracciu¹ ; T. Pulliam²

¹ *SLAC / INFN PADOVA*

² *Ohio State University*

The Event Reconstruction Control System of the BaBar experiment was redesigned in 2002, to satisfy the following major requirements: flexibility and scalability.

Because of its very nature, this system is continuously maintained to implement the changing policies, typical of a complex, distributed production environment.

In 2003, a major revolution in the BaBar computing model, the Computing Model 2, brought a particularly vast set of new requirements in various respects, many of which had to be discovered during the early production effort, and promptly dealt with. Particularly, the reconstruction pipeline was expanded with the addition of a third stage. The first fast calibration stage was kept running at SLAC, USA, while the two stages doing most of the computation were moved to the ~400 CPU reconstruction facility of INFN, Italy.

In this paper, we summarize the extent and nature of the evolution of the Control System, and we demonstrate how the modular, well engineered architecture of the system allowed to efficiently adapt and expand it, while making great reuse of existing code, leaving virtually intact the core layer, and exploiting the “engineering for flexibility” philosophy.

Plenary / 523

The impact of e-science

K. PEACH^{None}

Just as the development of the World Wide Web has had its greatest impact outside particle physics, so it will be with the development of the Grid.

E-science, of which the Grid is just a small part, is already making a big impact upon many scientific disciplines, and facilitating new scientific discoveries that would be difficult to achieve in any other way. Key to this is the definition and use of metadata.

Online Computing / 95

The introduction to BES computing environment

D. XU¹ ; G. CHEN¹

¹ COMPUTING CENTER, INSTITUTE OF HIGH ENERGY PHYSICS, CHINESE ACADEMY OF SCIENCES

BES is an experiment on Beijing Electron-Positron Collider (BEPC). BES computing environment consists of PC/Linux cluster and mainly relies on the free software. OpenPBS and Ganglia are used as job schedule and monitor system. With helps from CERN IT Division, CASTOR was implemented as storage management system. BEPC is being upgraded and luminosity will increase one hundred times comparing to current machine. The data produced by new BES-III detector will be about 700 Terabytes per year. To meet the computing demand, we proposed a solution based on PC/Linux/Cluster and SAN technology. CASTOR will be used to manage the storage resources of SAN. We started to develop a graphical interface for CASTOR. Some tests on data transmission performance of SAN environment were carried out. The result shows that I/O performance of SAN is better than that of traditional storage connection method including IDE, SCSI etc and it can satisfy BESIII experiment's demand for data processing.

Event Processing / 347

The new BaBar Analysis Model

D. Brown¹ ; D. Lange² ; E. Charles¹ ; G. De Nardo³ ; G. Finocchiaro⁴ ; L. Wilden⁵ ; P. Elmer⁶

¹ LAWRENCE BERKELEY NATIONAL LAB

² Lawrence Livermore National Lab

³ INFN Naples

⁴ INFN Frascati

⁵ NIKHEV Amsterdam

⁶ Princeton University

Auteur(s) contact: dave_brown@lbl.gov

This talk will describe the new analysis computing model deployed by BaBar over the past year. The new model was designed to better support the current and future needs of physicists analyzing data, and to improve BaBar's analysis computing efficiency.

The use of RootIO in the new model is described in other talks.

Babar's new analysis data content format contains both high and low level information, allowing physicists to pick a tradeoff between speed and precision/flexibility appropriate to their analysis.

The new format is customizable, allowing physicists to create analysis-specific content using simple and familiar tools.

Babar's new analysis processing model involves selecting events according to their physics content, and writing them together with analysis-customized content to dedicated output streams. Currently 120 such 'skims' are written as part of a periodic central processing cycle. Skims can be further reduced and customized as well as queried interactively in root-based applications. Skims and subskims retain links back to the original full event information. This processing model eliminates the need for large tuple production efforts by physicists.

The entire BaBar data sample is available in the new format, and the new model has been used to produce physics results presented at the summer HEP conferences.

We will also present reactions from the BaBar analysis community, and describe the issues that arose in deploying the new model.

Poster Session 2 - Tableau: 32 / 234

The role of legacy services within ATLAS DC2

DC2 production system team¹; J. Kennedy²

¹ ATLAS

² LMU Munich

Auteur(s) contact: john.kennedy@cern.ch

This paper presents an overview of the legacy interface provided for the ATLAS DC2 production system. The term legacy refers to any non-grid system which may be deployed for use within DC2. The reasoning behind providing such a service for DC2 is twofold in nature. Firstly, the legacy interface provides a backup solution should unforeseen problems occur while developing the grid based interfaces. Secondly, this system allows DC2 to use resources which have yet to deploy grid software, thus increasing the available computing power for the Data Challenge.

The aim of the legacy system is to provide a simple framework which is easily adaptable to any given computing system. Here the term computing system refers to the batch system provided at a given site and also to the structure of the computing and storage systems at that site. The legacy interface provides the same functionality as the grid based interfaces and is deployed transparently within the DC2 production system. Following the push-pull model implemented for DC2 the system pulls jobs from a production database and pushes them onto a given computing/batch system.

In a world which is becoming increasingly grid orientated this project allows us to evaluate the role of non-grid solutions in dedicated production environments. Experiences, both good and bad, gained during DC2 are presented and the future of such systems is discussed.

Plenary / 524

The role of scientific middleware in the future of HEP computing

M. Livny¹

¹ UNIVERSITY OF WISCONSIN

In the 18 months since the CHEP03 meeting in San Diego, the HEP community deployed the current generation of grid technologies in a veracity of settings. Legacy software as well as recently developed applications was interfaced with middleware tools to deliver end-to-end capabilities to HEP experiments in different stages of their life cycles. In a series of data challenges, reprocessing efforts and data distribution activities the community demonstrated the benefits distributed computing can offer and the power a range of middleware tools can deliver. After running millions of jobs, moving tera-bytes of data, creating millions of files and resolving hundreds of bug reports, the community also exposed the limitations of these middleware tools. As we move to the next level of challenges, requirements and expectations, we must also examine the methods and procedures we employ to develop, implement and maintain our common suite of middleware tools. The talk will focus on the role common middleware developed by the scientific community can and should play in the software stack of current and future HEP experiments.

Event Processing / 67

The simulation for the ATLAS experiment: present status and outlook

A. DELL'ACQUA¹ ; A. NAIRZ¹ ; A. Rimoldi² ; D. COSTANZO³ ; J. BOUDREAU⁴ ; M. GALLAS⁵ ; V. TSULAIA⁴

¹ CERN, Geneva, Switzerland

² PAVIA UNIVERSITY & INFN

³ LBL, USA

⁴ University of Pittsburgh, USA

⁵ CERN

Auteur(s) contact: adele.rimoldi@cern.ch

The simulation for the ATLAS experiment is presently operational in a full OO environment and it is presented here in terms of successful solutions to problems dealing with application in a wide community using a common framework. The ATLAS experiment is the perfect scenario where to test all applications able to satisfy the different needs of a big community. Following a well stated strategy of transition from the GEANT3 to the GEANT4-based simulation, a good validation programme during the last months confirmed the characteristics of reliability, performance and robustness of this new tool in comparison with the results of the previous simulation. Generation, simulation and digitization steps on different full sets of physics events were tested in terms of performance and robustness in comparisons with the same samples undergoing the old GEANT3-based simulation. The simulation program is simultaneously tested on all different testbeam setups characterizing the R&D programme of all subsystems belonging to the ATLAS detector with comparison to real data in order to validate the physics content and the reliability in the detector description of each component.

Computer Fabrics / 464

The status of Fermilab Enstore Data Storage System

A. Moibenko¹ ; C-H. Huang¹ ; D. Petravick¹ ; E. Berman¹ ; J. Bakken¹ ; M. Zalokar¹

¹ *FERMI NATIONAL ACCELERATOR LABORATORY, USA*

Auteur(s) contact: moibenko@fnal.gov

Fermilab has developed and successively uses Enstore Data Storage System. It is a primary data store for the Run II Collider Experiments, as well as for the others. It provides data storage in robotic tape libraries according to requirements of the experiments. High fault tolerance and availability, as well as multilevel priority based request processing allows experiments to effectively store and access data stored in the Enstore, including storing raw data from data acquisition systems. The distributed structure and modularity of Enstore allow to scale the system and add more storage equipment as the requirements grow. Currently Fermilab Data Storage System storage system Enstore includes 5 robotic tape libraries, 96 tape drives of different type. Amount of data stored in the system is ~1.7 PetaBytes. Users access Enstore directly using a special command. They also can use ftp, grid ftp, SRM interfaces to dCache system, that uses Enstore as its lower layer storage.

Distributed Computing Systems and Experiences / 136

Tier-1 and Tier-2 Real-time Analysis experience in CMS Data Challenge 04

A. Fanfani¹ ; A. Pierro² ; C. Grandi¹ ; D. Bonacorsi³ ; F. Ambroglini⁴ ; F. Fanzago⁵ ; G. Donvito⁶ ; J. Hernandez⁷ ; L. Silvestris² ; M. Corvo⁵ ; N. De Filippis⁸ ; V. Innocente⁹

¹ *INFN Bologna*

² *INFN Bari*

³ *CNAF-INFN Bologna*

⁴ *Universita' di Perugia and INFN*

⁵ *INFN Padova*

⁶ *Universita' di Bari*

⁷ *CIEMAT Madrid*

⁸ *UNIVERSITA' DEGLI STUDI DI BARI AND INFN*

⁹ *CERN, PH/SFT*

During the CMS Data Challenge 2004 a realtime analysis was attempted at INFN and PIC Tier-1 and Tier-2s in order to test the ability of the instrumented methods to quickly process the data.

Several agents and automatic procedures were implemented to perform the analysis at the Tier-1/2 synchronously with the data transfer from Tier-0 at CERN. The system was implemented in the Grid LCG-2 environment and allowed on-the-fly job preparation and subsequent submission to the Resource Broker as new data come along. Running job accessed data from the Storage Elements through POOL via remote file protocol, whenever possible, or copying them locally with gridftp.

Job monitoring and bookkeeping was performed using BOSS. Details of the procedures adopted to run the analysis jobs and the expected results are described.

An evaluation of the ability of the system to maintain an analysis rate at Tier-1 and Tier-2 comparable with the data transfer rate is

also presented.

The results on the analysis timeline, the statistics of submitted jobs, the overall efficiency of the GRID services and the overhead introduced by the agents/procedures are reported. Performances and possible bottlenecks of the whole procedure are discussed.

Poster Session 2 - Tableau: 13 / 335

Tools for GRID deployment of CDF offline and SAM data handling systems for Summer 2004 computing.

A. Baranovski¹ ; A. Kreymer² ; A. Lyon³ ; A. Sill⁴ ; F. Ratnikov⁵ ; G. Garzoglio³ ; I. Terekhov¹ ; J. Trumbo¹ ; L. Loebel Carpenter¹ ; M. Burgon-Lyon⁶ ; M. Leslie⁷ ; R. Herber³ ; R. Illingworth³ ; R. Kennedy³ ; R. St.Denis⁶ ; S. Belforte⁸ ; S. Stonjek¹ ; S. Veseli¹ ; S. White¹ ; U. Kerzel⁹ ; V. Bartsch⁷ ; W. Merritt³

¹ *Fermilab*

² *FERMILAB*

³ *FERMI NATIONAL ACCELERATOR LABORATORY*

⁴ *Texas Tech. University*

⁵ *Rutgers University*

⁶ *Glasgow University*

⁷ *Oxford University*

⁸ *INFN*

⁹ *Karlsruhe University*

Auteur(s) contact: kreymer@fnal.gov

The Fermilab CDF Run-II experiment is now providing official support for remote computing, expanding this to about 1/4 of the total CDF computing during the Summer of 2004.

I will discuss in detail the extensions to CDF software distribution and configuration tools and procedures, in support of CDF GRID/DCAF computing for Summer 2004. We face the challenge of unreliable networks, time differences, and remote managers with little experience with this particular software.

We have made the first deployment of the SAM data handling system outside its original home in the D0 experiment.

We have deployed to about 20 remote CDF sites.

We have created light weight testing and monitoring tools to assure that these sites are in fact functional when installed.

We are distributing and configuring both client code within CDF code releases, and the SAM servers to which the clients connect.

Procedures which once took days are now performed in minutes.

These tools can be used to install SAM servers for D0 and other experiments.

Networks permitting, we will give a live SAM installation demonstration.

We have separated the data handling components from the main CDF offline code releases by means of shared libraries, permitting live upgrades to otherwise frozen code.

We now use a special 'development lite' release to ensure that all sites have the latest tools available.

We have put substantial effort into revision control, so that essentially all active CDF sites are running exactly the same code.

Poster Session 2 - Tableau: 16 / 78

Toward a Grid Technology Independent Programming Interface for HEP Applications

A. Delgado Peris¹ ; A. Sciaba¹ ; F. Donno¹ ; P. Mendez Lorenzo¹ ; R. Santinelli¹ ; S. Campana¹

¹ CERN

Auteur(s) contact: flavia.donno@cern.ch

In the High Energy Physics (HEP) community, Grid technologies have been accepted as solutions to the distributed computing problem. Several Grid projects have provided software in the last years. Among of all them, the LCG - especially aimed at HEP applications - provides a set of services and respective client interfaces, both in the form of command line tools as well as programming language APIs in C, C++, Java, etc.

Unfortunately, the programming interface presented to the end user (the physicist) is often not uniform or provides different levels of abstractions. In addition, Grid technologies face a constant change and an improvement process and it is of major importance to shield changes of underlying technology to the end users. As services evolve and new ones are introduced, the way users interact with them also changes.

These new interfaces are often designed to work at a different level and with a different focus than the original ones. This makes it hard for the end user to build Grid applications.

We have analyzed the existing LCG programming environment and

identified several ways to provide high-level technology independent interfaces. In this article, we describe the use cases we were presented by the LCG experiments and the specific problems we encountered in documenting existing APIs and providing usage examples. As a main contribution, we also propose a prototype high-level interface for the information, authentication and authorization systems that is now under test on the LCG EIS testbed by the LHC experiments.

Event Processing / 435

Track reconstruction in high density environment

I. Belikov¹ ; K. Safarik¹ ; M. Ivanov¹ ; P. Hristov¹ ; T. Kuhr¹

¹ CERN

Auteur(s) contact: marian.ivanov@cern.ch

Tracks finding and fitting algorithm in ALICE Time projection chamber (TPC) and Inner Tracking System (ITS) based on the Kalman-filtering are presented. The filtering algorithm is able to cope with non-Gaussian noise and ambiguous measurements in high-density environments. The tracking algorithm consists of two parts: one for the TPC and one for the prolongation into the ITS. The occupancy in the TPC can reach up to 40 %.

Usually, due to the overlaps, a number of points along the track are lost or significantly displaced. At first the clusters are found and the space points are reconstructed. The shape of a cluster provides information about the overlap. An unfolding algorithm is applied for points with distorted shapes. Then, the expected space point error is estimated using information about the cluster shape and track parameters. Further, the available information about local track overlap is used. In the TPC-ITS matching, the distance between the TPC and the ITS sensitive volume is rather large and the track density inside the ITS is so high that the straightforward continuation of the tracking procedure is ineffective. Using only chi2 minimisation there is a high probability of assigning a wrong hit to the track. Therefore for each TPC track a candidate tree of the possible track prolongation in the ITS is build. Finally the most probable track candidates are chosen. The approach have been implemented within the ALICE simulation/reconstruction framework (ALIROOT), and algorithm's efficiency have been estimated using the ALIROOT Monte Carlo data.

Poster Session 3 - Tableau: 17 / 70

Tracking of long lived hyperons in silicon detector at CDF.

D. Litvintsev¹ ; E. Gerchtein²

¹ *FNAL*

² *CMU*

Auteur(s) contact: lena@fnal.gov

Long lived charged hyperon, Ξ and Ω , are capable of travelling significant distances producing hits in the silicon detector, before decaying into $\Lambda^0\pi$ and Λ^0K pairs, respectively. This gives unique opportunity of reconstructiong hyperon tracks. We have developed a dedicated "outside-in" tracking algorithm that is seeded by 4-momentum and decay vertex of the long lived hyperon reconstructed by its decay products. The tracking of hyperons in the silicon detector results in a dramatic reduction of the combinatorial background and an improvement of the momentum resolution compared with the standard reconstruction using final decay products.

Using a super clean sample of Ξ hyperons CDF observed charmed-strange baryon isodoublet Ξ_c^0 and Ξ_c^+ for the first time in $p\bar{p}$ collision. Ξ hyperons were used for the search for exotic $S = -2$ baryons decaying into $\Xi\pi$.

Poster Session 3 - Tableau: 30 / 120

Transparently managing time varying conditions and detector data on ATLAS.

C. Leggett¹

¹ *LAWRENCE BERKELEY NATIONAL LABORATORY*

Auteur(s) contact: charles.g.leggett@gmail.com

It is essential to provide users transparent access to time varying data, such as detector misalignments, calibration parameters and the like. This data should be automatically updated, without user intervention, whenever it changes. Furthermore, the user should be able to be notified whenever a particular datum is updated, so as to perform actions such as re-caching of compound results, or performing computationally intensive task only when necessary. The user should only have to select a particular calibration scheme or time interval, without having to worry about explicitly updating data on an event by event basis. In order to minimize database activity, it is important that the system only manage the parameters that are actively used in a particular job, making updates only on demand. For certain situations however, such as testbeam environments, pre-caching of data is essential, so the system must also be able to pre-load all relevant data at the start of a run, and avoid further updates to the data.

In this talk we present the scheme for managing time varying data and their associated intervals of validity, as used in the Athena framework on ATLAS, which features automatic updating of conditions data occurring invisibly to the user; automatic and explicit registration of objects of interest; callback function hierarchies; and abstract conditions database interfaces.

Event Processing / 280

Update On the Status of the FLUKA Monte Carlo Transport Code

A. Empl¹ ; A. Fasso² ; A. Ferrari³ ; A. Ottolenghi⁴ ; D. Scannicchio⁴ ; E. Gadioli⁵ ; F. Ballarini⁴ ; F. Cerutti⁵ ; G. Battistoni⁵ ; G. SMIRNOV⁶ ; J. Ranft⁷ ; K. LEE⁶ ; K. TSOULOU⁸ ; L. Pinsky⁹ ; M V. Garzelli¹⁰ ; M. CAMPANELLA¹¹ ; M. CARBONI¹² ; M. PELLICIONI¹² ; N. Zapp¹ ; P. Sala¹³ ; S. MURARO¹¹ ; S. Roesler¹⁴ ; T. RANCATI¹¹ ; V. ANDERSON⁶ ; V. PARINI¹⁵ ; V. Vlachoudis¹⁴

¹ *University of Houston*

² *SLAC*

³ *CERN & Universita' and INFN Milano*

⁴ *Universita' and INFN Pavia*

⁵ *Universita' and INFN Milano*

⁶ *University of Houston, USA*

⁷ *University of Leipzig*

⁸ *CERN, Geneva, Switzerland*

⁹ *UNIVERSITY OF HOUSTON*

¹⁰ *Univeristy and INFN, Milano, Italy*

¹¹ *INFN Milano, Italy*

¹² *INFN Frascati, Italy*

¹³ *ETH Zurich & INFN Milan*

¹⁴ *CERN*

¹⁵ *INFN Pavia, Italy*

Auteur(s) contact: pinsky@uh.edu

The FLUKA Monte Carlo transport code is a well-known simulation tool in High Energy Physics. FLUKA is a dynamic tool in the sense that it is being continually updated

and improved by the authors. Here we review the progresses achieved in the last year on the physics models. From the point of view of hadronic physics, most of the effort is still in the field of nucleus–nucleus interactions. The currently available version of FLUKA already includes the internal capability to simulate inelastic nuclear interactions beginning with lab kinetic energies of 100 MeV/A up to the highest accessible energies by means of the DPMJET-II.5 event generator to handle the interactions for >5 GeV/A and rQMD for energies below that. The new developments concern, at high energy, the embedding of the DPMJET-III generator, which represent a major change with respect to the DPMJET-II structure. This will also allow to achieve a better consistency between the nucleus-nucleus section with the original FLUKA model for hadron-nucleus collisions. Work is also in progress to implement a third event generator model based on the Master Boltzmann Equation approach, in order to extend the energy capability from 100 MeV/A down to the threshold for these reactions. In addition to these extended physics capabilities, structural changes to the programs input and scoring capabilities are continually being upgraded. In particular we want to mention the upgrades in the geometry packages, now capable of reaching higher levels of abstraction. Work is also proceeding to provide direct import into ROOT of the FLUKA output files for analysis and to deploy a user-friendly GUI input interface.

Poster Session 2 - Tableau: 17 / 431

Usage of ALICE Grid middleware for medical applications

E. Lopez Torres¹ ; F. Carminati² ; F. Rademakers² ; P. BUNCIC² ; P. Cerello³ ; P. Hristov² ; S. Bagnasco³ ; S. Cheran³

¹ *CEADEN - Cuba*

² *CERN*

³ *INFN Torino*

Auteur(s) contact: cerello@to.infn.it, federico.carminati@cern.ch

Breast cancer screening programs require managing and accessing a huge amount of data, intrinsically distributed, as they are collected in different Hospitals. The development of an application based on Computer Assisted Detection algorithms for the analysis of digitised mammograms in a distributed environment is a typical GRID use case. In particular, AliEn (ALICE Environment) services, whose development was carried on by the ALICE Collaboration, were used to configure a dedicated Virtual Organisation; a PERL-based interface to AliEn commands allows the registration of new patients and mammograms in the AliEn Data Catalogue as well as queries to retrieve images associated to selected patients. The analysis of selected mammograms can be performed interactively, making use of PROOF services, or taking advantage of the AliEn capabilities to generate “sub-jobs”; each of them analyzes the fraction of the selected sample stored on a site, and the results are merged. All the required functionality is available: by the end of 2004 a working prototype is foreseen, with an AliEn Client installed in each of the Hospitals participating to the INFN-funded MAGIC-5 project. The same approach will be applied in the near future in two other application areas:

- Lung cancer screening, equivalent to the mammographic screening from the middleware point of view, where Computer Assisted Detection algorithms are being developed;
- Diagnosis of the Alzheimer disease, where the application is intrinsically distributed: it should, in fact, compare the PET-generated image to a set of reference images which are scattered on many sites and merge the results.

Poster Session 2 - Tableau: 34 / 102

Usage statistics and usage patterns on the NorduGrid: Analyzing the logging information collected on one of the largest production Grids of the world.

A. KONSTANTINOV¹ ; A. WAANANEN² ; B. Konya³ ; F. OULD-SAAD⁴ ; J. R. HANSEN² ; J.L. Nielsen² ; K. PAJCHEL⁴ ; M. Ellert⁵ ; O. SMIRNOVA⁶ ; P. Eerola⁶ ; T. EKELOF⁷ ; U. ERKARSLAN⁴

¹ *Vilnius University, Lithuania*

² *Niels Bohr Institute, Copenhagen, Denmark*

³ *LUND UNIVERSITY, NORDUGRID PROJECT*

⁴ *University of Oslo, Norway*

⁵ *CERN*

⁶ *Lund University, Sweden*

⁷ *Uppsala University, Sweden*

Auteur(s) contact: balazs.konya@hep.lu.se

The Nordic Grid facility (Nordugrid) came into production operation during the summer of 2002 when the Scandinavian Atlas HEP group started to use the Grid for the Atlas Data Challenges and was thus the first Grid ever contributing to an Atlas production. Since then, the Grid facility has been in continuous 24/7 operation offering an increasing number of resources to a growing set of active users coming from various scientific areas including chemistry, biology, informatics. As of today the Grid has grown into one of the largest production Grids of the world continuously running Grid jobs on the more than 30 Grid-connected sites which offer over 2000 CPUs.

This article will start with a short overview of the design and implementation of the Advanced Resource Connector (ARC), the Nordugrid middleware, which delivers reliable Grid services to the Nordugrid production facility. This will be followed by a presentation of the logging facility of Nordugrid, describing the logging service and the collected information. The main part of the talk will focus on the analysis of the collected logging information: usage statistics, usage patterns (what is a typical grid job on the Nordugrid looks like?). Use cases from different application domains will also be discussed.

References:

-Nordugrid live: www.nordugrid.org -> Grid Monitor

-Atlas Data-Challenge 1 on Nordugrid: <http://arxiv.org/abs/physics/0306013>

Distributed Computing Systems and Experiences / 307

Use of Condor and GLOW for CMS Simulation Production

D. Bradley¹ ; M. Livny¹ ; S. Dasu¹ ; S. Rader¹ ; V. Puttabuddhi¹ ; W. Smith¹

¹ *UNIVERSITY OF WISCONSIN*

Auteur(s) contact: dasu@hep.wisc.edu

The University of Wisconsin distributed computing research groups developed a software system called Condor for high throughput computing using commodity hardware. An adaptation of this software, Condor-G, is

part of Globus grid computing toolkit. However, original Condor has additional features that allows building of an enterprise level grid. Several UW departments have Condor computing pools that are integrated in such a way as to flock jobs from one pool to another as resources become available. An interdisciplinary team of UW researchers recently built a new distributed computing facility, the Grid Laboratory of Wisconsin (GLOW). In total Condor pools in the UW have about 2000 Intel CPUs (P-III and Xeon) which are available for scientific computation. By exploiting special features of Condor such as checkpointing and remote IO we have generated over 10 million fully simulated CMS events. We were able to harness about 260 CPU-days per day for a period of 2 months when we were operational late fall. We have scaled to using 500 CPUs concurrently when opportunity to exploit unused resources in laboratories on our campus. We have built a scalable job submission and tracking system called Jug using Python and mySQL which enabled us to scale to run hundreds of jobs simultaneously. Jug also ensured that the data generated is transferred to US Tier-I center at Fermilab. We have also built a portal to our resources and participated in Grid2003 project. We are currently adapting our environment for providing analysis resources. In this paper we will discuss our experience and observations regarding the use of opportunistic resources, and generalize them to wider grid computing context.

Poster Session 1 - Tableau: 31 / 394

Using HEP Systems to Provide Storage for Biologists

M. Calef¹

¹ VANDERBILT UNIVERSITY

Auteur(s) contact: matthew.t.calef@vanderbilt.edu

Protein analysis, imaging, and DNA sequencing are some of the branches of biology where growth has been enabled by the availability of computational resources. With this growth, biologists face an associated need for reliable, flexible storage systems. For decades the HEP community has been driving the development of such storage systems to meet their own needs. Two of these systems - the dCache disk caching system and the Enstore hierarchical storage manager - are viable candidates for addressing the storage needs of biologists. Both incorporate considerable experience from the HEP community.

While biologists have much to gain from the HEP community's experience with storage systems, they face several issues that are unique to the biological sciences. There is a wider diversity in experiments, in number and size of datafiles, and in client operating systems in biology than there is in HEP. Patient information must be kept confidential. Disparate IT departments set up firewalls that separate client systems and the storage system.

Vanderbilt University is developing a storage system with the goal of meeting biologists' needs. This system will use Enstore for its robustness and reliability, and will use the flexible door-based architecture of dCache to provide storage services to biologists via web-portal, the dCache copy command, and custom applications. This system will be deployed using an automated tape library, several secure central servers, and nodes placed near biologists' existing compute infrastructure to ensure locality of caches and secure data channels between researchers and the central servers.

Grid Security / 164

Using Nagios for intrusion detection

E. PEREZ-CALLE¹ ; F.J. Rodriguez Calonge¹ ; M. Cardenas Montes¹¹ CIEMAT**Auteur(s) contact:** miguel.cardenas@ciemat.es

Implementing strategies for secured access to widely accessible clusters is a basic requirement of these services, in particular if GRID integration is sought for. This issue has two complementary lines to be considered: security perimeter and intrusion detection systems. In this paper we address aspects of the second one.

Compared to classical intrusion detection mechanisms, close monitoring of computer services can substantially help to detect intrusion signs. Having alarms indicating the presence of an intrusion into the system, allows system administrators to take fast actions to minimize damages and stop diffusion towards other critical systems.

One possible monitoring tool is Nagios (www.nagios.org), a powerful GNU tool with capacity to observe and collect information about a variety of services, and trigger alerts.

In this paper we present the work done at CIEMAT, where we have applied these directives to our local cluster. We have implemented a system to monitor the hardware and system sensitive information. We describe the process and show through different simulated security threads how does our implementation respond to it.

Poster Session 2 - Tableau: 45 / 163

Using Tripwire to check cluster system integrity

E. Perez-Calle¹ ; F.J. Rodriguez Calonge¹ ; M. Cardenas Montes¹¹ CIEMAT**Auteur(s) contact:** elio.perez@ciemat.es

Expansion of large computing fabrics/clusters throughout the world would create a need for stricter security. Otherwise any system could suffer damages such as data loss, data falsification or misuse.

Perimeter security and intrusion detection system (IDS) are the two main aspects that must be taken into account in order to achieve system security.

The main target of an intrusion detection system is early detection in the previously mentioned cases, as a way to minimize any damage in data contained in the system.

Tripwire is one of the most powerful IDSs and is widely used as a security tool by the community of network administrators. Tripwire is oriented to monitor the status of files and directories, being able to detect the lightest change suffered by them.

At Ciemat, Tripwire has been used to monitor our local clusters, involved in GRID projects such as implementation of LCG prototypes, to guarantee the integrability of data generated, and stored there. It is

used as well to monitor any modification of operating system files and any other scientific core software.

Event Processing / 77

Using the reconstruction software, ORCA, in the CMS datachallenge

S. Wynhoff¹

¹ *PRINCETON UNIVERSITY*

Auteur(s) contact: stephan.wynhoff@cern.ch

We report on the software for Object-oriented Reconstruction for CMS Analysis, ORCA. It is based on the Coherent Object-oriented Base for Reconstruction, Analysis and simulation (COBRA) and used for digitization and reconstruction of simulated Monte-Carlo events as well as testbeam data.

For the 2004 data challenge the functionality of the software has been extended to store collections of reconstructed objects (DST) as well as the previously storable quantities (Digis) in multiple, parallel streams.

We describe the structure of the DST, the way to ensure and store the configuration of reconstruction algorithms that fill the collections of reconstructed objects as well as the relations between them. Also the handling of multiple streams to store parts of selected events is discussed. The experience from the implementation used early 2004 and the modifications for future optimization of reconstruction and analysis are presented.

Poster Session 3 - Tableau: 31 / 244

Validation of the GEANT4 Bertini Cascade model and data analysis using the Parallel ROOT Facility on a Linux cluster

A. Heikkinen¹ ; T. Linden²

¹ *Helsinki Institute of Physics HIP*

² *HELSINKI INSTITUTE OF PHYSICS*

Auteur(s) contact: tomas.linden@helsinki.fi

Validation of hadronic physics processes of the Geant4 simulation toolkit is a very important task to ensure adequate physics results for the experiments being built at the Large Hadron Collider. We report on simulation results obtained using the Geant4 Bertini cascade double-differential production cross-sections for various target materials and incident hadron kinetic energies between 0.1-10 GeV [1, 2].

The cross-section benchmark study in this work has been performed using a Linux cluster set up with the Red Hat Linux based NPACI Rocks Cluster Distribution. For analysis of the validation data we have used the Parallel ROOT Facility (PROOF). PROOF has been designed for setting up a

parallel data analysis environment in an inhomogeneous computing environment. Here we use a homogeneous Rocks cluster and automatic class generation for PROOF event data-analysis [3].

[1] J. Beringer, “(p, xn) Production Cross Sections: A benchmark Study for the Validation of Hadronic Physics Simulation at LHC”, CERN-LCGAPP-2003-18.

[2] A. Heikkinen, N. Stepanov, and J.P. Wellisch, “Bertini intra-nuclear cascade implementation in Geant4”, arXiv: nucl-th/0306008

[3] F. Rademakers, M. Goto, P. Canal, R. Brun, “ROOT Status and Future Developments”, arXiv: cs.SE/0306078

Event Processing / 106

Vertex finding and B-tagging algorithms for the ATLAS Inner Detector

A. Wildauer¹ ; F. Akesson²

¹ UNIVERSITY OF INNSBRUCK

² CERN

Auteur(s) contact: andreas.wildauer@cern.ch

For physics analysis in ATLAS, reliable vertex finding and fitting algorithms are important. In the harsh environment of the LHC (~ 23 inelastic collisions every 25 ns) this task turns out to be particularly challenging. One of the guiding principles in developing the vertexing packages is a strong focus on modularity and defined interfaces using the advantages of object oriented C++. The benefit is the easy expandability of the vertexing with additional fitting strategies integrated in the Athena framework.

Various implementations of algorithms and strategies dedicated to primary and secondary vertex reconstruction using the full reconstruction of simulated ATLAS events are presented.

Primary and secondary vertex finding is essential for the identification of b-jets in a reconstructed event. Results from a modular and expandable b-tagging algorithm are shown using the presented strategies for vertexing.

Poster Session 3 - Tableau: 27 / 404

Visualisation of the ROOT geometries (TVolume and TGeoVolume) with Coin-based model.

A. Kulikov¹ ; D. Turkin¹ ; M. Panebrattsev¹ ; V. Fine²

¹ Joint Institute for Nuclear Research, Dubna, Russia.

² Brookhaven National Laboratory, Upton, USA

Auteur(s) contact: shurik@sunhe.jinr.ru

Using the modern 3D visualization software and hardware to represent the object models of the HEP detectors would create the impressive pictures of events and the detail views of the detectors facilitating the design, simulation and data analysis and representation the huge amount of the information flooding the modern HEP experiments. In this paper we represent the work made by members of STAR collaboration from Laboratory of High Energy Physics JINR Dubna. This work devoted to visualisation of the STAR detector geometry. Initially the detector geometry is described by means of specific AGE - geometry specification language and it can be converted to either TVolume or TGeoVolume type object of ROOT environment using of specially developed software. We created class library for conversion of the ROOT OO model of the detector from ROOT environment to the text "iv" file. Our class library assumes the conversion of ROOT OO models to Coin-based C++ representation and the Coin-based 3d Viewer with cutting / highlighting / selecting pieces of the image features. Since the class library implementation is free of the STAR experiment specific it can be used to visualize any detector geometry represented in the ROOT environment. The results of our work can be downloaded from the LHE web server.

Poster Session 3 - Tableau: 58 / 246

Volume-based representation o the magnetic field

N. AMAPANE¹ ; T. Todorov² ; V. Andreev³ ; V. Drollinger¹ ; V. Karimaki⁴ ; V. Klyukhin⁵

¹ INFN

² CERN/IReS

³ UCLA

⁴ HELSINKI-HIP

⁵ MSU/CERN

Auteur(s) contact: theodore.todorov@cern.ch

The simulation, reconstruction and analysis software access to the magnetic field has large impact both on CPU performance and on accuracy.

An approach based on a volume geometry is described. The volumes are constructed in such a way that their boundaries correspond to field discontinuities, which are due to changes in magnetic permeability of the materials. The field in each volume is contiguous.

The field in each volume is interpolated from a regular grid of values resulting from a TOSCA calculation. In case a parameterization is available for some volumes it is used instead of the grid interpolation.

Global access to the magnetic field values requires efficient search for the volume that contains a global point. An algorithm that exploits explicitly the layout and the symmetries of the detector is presented.

The main clients of the magnetic field, which are the simulation (geant4) and propagation of track parameters and errors in the reconstruction, can be made aware of the magnetic field volumes by connecting the per-volume magnetic field providers in their respective geometries to the corresponding volume in the magnetic

field geometry. In this way the global volume search is by-passed and the access to the field is sped up significantly.

Poster Session 1 - Tableau: 2 / 369

WAN Emulation Development and Testing at Fermilab

A. Bobyshev¹ ; M. Ernst² ; P. Demar¹ ; R. Rechenmacher¹

¹ *FERMILAB*

² *DESY*

Auteur(s) contact: bobyshev@fnal.gov

The Compact Muon Solenoid (CMS) experiment at CERN's Large Hadron Collider (LHC) is scheduled to come on-line in 2007. Fermilab will act as the CMS Tier-1 center for the US and make experiment data available to more than 400 researchers in the US participating in the CMS experiment. The US CMS Users Facility group, based at Fermilab, has initiated a project to develop a model for optimizing movement of CMS experiment data between CERN and the various tiers of US CMS data centers. Fermilab has initiated a project to design a WAN emulation facility which will enable controlled testing of unmodified or modified CMS applications and TCP implementations locally under conditions that emulate WAN connectivity. The WAN emulator facility is configurable for latency, jitter, and packet loss. The initial implementation is based on the NISTnet software product. In this paper we will describe the status of this project to date, the results of validation and comparison of performance measurements obtained in emulated and real environment for different applications including multistreams GridFTP. We also will introduce future short term and intermediate term plans, as well as outstanding problems and issues.

Event Processing / 395

WIRED 4 - A generic Event Display plugin for JAS 3

M. Donszelmann¹

¹ *SLAC*

Auteur(s) contact: mark.donszelmann@slac.stanford.edu

WIRED 4 is a experiment independent event display plugin module for JAS 3 (Java Analysis Studio) generic analysis framework. Both WIRED and JAS are written in Java.

WIRED, which uses HepRep (HEP Representables for Event Display) as its input format, supports viewing of events using either conventional 3D projections as well as specialized projections such as a fish-eye or a rho-Z projection. Projections allow the user to scale, rotate, position or change parameters on the plot as he wishes. All interactions are handled as separate edits which can be undone and/or redone, so the user can try out things and get back to a previous situation. All edits are scriptable by any of the scripting languages supported by JAS, such as pnuts, jython or java. Hits and tracks can be picked to display physics information and cuts can be made on physics parameters to allow the user to filter the number of objects drawn into the plot. Multiple event display plots can be laid out on pages combined with histograms and other plots,

available from JAS itself or from other plugin modules.
Configuration information on the state of all plots can be saved and restored allowing the user to save his session, share it with others or later continue where he left off.

This version of WIRED is written to be easily extensible by the user/developer. Projections, representations, interaction handlers and edits are all services and new ones can be added by writing additional plugins.
Both JAS 3 and WIRED 4 are built on top of the FreeHEP Java Libraries, which support a multitude of vector graphics output formats, such as PostScript, PDF, SVG, SWF and EMF, allowing document quality output of event display plots and histograms.

References:

<http://wired.freehep.org>
<http://jas.freehep.org/jas3>
<http://java.freehep.org>

Plenary / 520

Welcome Address

W. VON RUEDEN^{None}

Wide Area Networking / 348

Wide Area Network Monitoring system for HEP experiments at Fermilab

C. Logg¹ ; L. Cottrell¹ ; M. Grigoriev²

¹ SLAC, USA

² FERMILAB, USA

Auteur(s) contact: maxim@fnal.gov

Large, distributed HEP collaborations, such as D0, CDF and US-CMS, depend on stable and robust network paths between major world research centers. The evolving emphasis on data and compute Grids increases the reliance on network performance.
FermiLab's experimental groups and network support personnel identified a critical need for WAN monitoring to ensure the quality and efficient utilization of such network paths. This has led to the development of the Network Monitoring system we will present in this paper.
The system evolved from the IEPM-BW project, started at SLAC two years ago.
At Fermilab it has developed into a fully functional infrastructure with bi-directional active network probes and path characterizations. It is based on the Iperf achievable throughput tool, Ping and Synack to test ICMP/TCP connectivity, Pipechar and Traceroute to test, compare and report hop-by-hop network path characterization, and real file transfer performance by BBFTP and GridFTP. The Monitoring system has an extensive web-interface and all the data is available through standalone SOAP web services or by a MonALISA client.
Also in this paper we will present a case study of network path

asymmetry and abnormal performance between FNAL and SDSC which was discovered and resolved by utilizing the Network Monitoring system.

Event Processing / 402

Writing Extension Modules (Plug-ins) for JAS3

Auteur(s): T. Johnson¹

Co-auteur(s) Mark Donszelmann¹ ; Massimiliano Turri¹ ; Victor Serbo¹

¹ SLAC

Auteur(s) contact: tonyj@slac.stanford.edu

JAS3 is a general purpose, experiment independent, open-source, data analysis tool. JAS3 includes a variety of features, including histogramming, plotting, fitting, data access, tuple analysis, spreadsheet and event display capabilities. More complex analysis can be performed using several scripting languages (pnuts, jython, etc.), or by writing Java analysis classes. All of these features are provided by loosely coupled “plug-in” modules which are installed into the JAS3 base application framework.

In this presentation we will describe the JAS3 plug-in architecture, and explain how different plug-ins can interact via service interfaces and event dispatch mechanisms. We will demonstrate how this architecture makes it possible for individual plug-ins to be added, removed or upgraded to customize the application. We will then give an overview of how to design new experiment or domain specific plug-ins to extend the functionality of JAS3 for your own requirements, or to provide general purpose components for use by others.

Core Software / 412

XML I/O in ROOT

H. Essel¹ ; R. Brun² ; S. Linev³

¹ GSI, Darmstadt

² CERN

³ GSI

Auteur(s) contact: s.linev@gsi.de, h.essel@gsi.de

Till now, ROOT objects can be stored only in a binary ROOT specific file format. Without the ROOT environment the data stored in such files are not directly accessible. Storing objects in XML format makes it easy to view and edit (with some restriction) the object data directly. It is also plausible to use XML as exchange format with other applications. Therefore XML streaming has been implemented in ROOT. Any object which is in the ROOT dictionary can be stored/retrieved in XML format. Two layouts of object representation in XML are supported: class-dependent and generic. In the first case all XML tag names are derived from class and member names. To avoid name intersections, XML namespaces for each class are used. A Document Type Definition (DTD) file is automatically generated for each class (or set of classes). It can be used to validate the structure of the XML document. The generic layout of XML files includes tag names like “Object”, “Member”, “Item” and so on. In this case the DTD is common for all produced XML files. Further development is required to provide tools for accessing created XML files from other

applications like: pure C++ code without ROOT libraries and dictionaries, Java and so on.

Poster Session 2 - Tableau: 24 / 128

XTNetFile, a fault tolerant extension of ROOT TNetFile

A. Dorigo¹ ; A. Hanushevsky² ; F. Furano³ ; P. Elmer⁴

¹ *INFN PADOVA*

² *SLAC*

³ *INFN Padova*

⁴ *Princeton University*

Auteur(s) contact: alvise.dorigo@pd.infn.it

This paper describes XTNetFile, the client side of a project conceived to address the high demand data access needs of modern physics experiments such as BaBar using the ROOT framework. In this context, a highly scalable and fault tolerant client/server architecture for data access has been designed and deployed which allows thousands of batch jobs and interactive sessions to effectively access the data repositories basing on the XROOTD data server, a complex extension of the rootd daemon. The majority of the communication problems are handled by the design of the client/server mechanism and the communication protocol.

This allows us to build distributed data access systems which are highly robust, load balanced and scalable to an extent which allows 'no jobs to fail'.

Furthermore XTNetFile ensures backward compatibility with the 'old' rootd server by using same API as the existing ROOT TFile/TNetFile classes.

The code is designed with a high degree of modularity that allows to build other interfaces, such as administrative tools, based on the same communication layer. In addition the client plugin can also be used to read other types of (non-ROOT I/O) data files, providing the same benefits.

Computer Fabrics / 233

dCache, LCG Storage Element and enhanced use cases

P. Fuhrmann¹

¹ *DESY*

Auteur(s) contact: patrick.fuhrmann@desy.de

The dCache software system has been designed to manage a huge amount of individual disk storage nodes and let them appear under a single file system root. Beside a variety of other features, it supports the GridFtp dialect, implements the Storage Resource Manager interface (SRM V1) and can be linked against the CERN GFAL software layer. These abilities makes dCache a perfect Storage Element in the context of LCG and

possibly future grid initiatives as well.

During the last year, dCache has been deployed at dozens of Tier-I and Tier-II centers for the CMS and CDF experiments in the US and Europe, including Fermilab, Brookhaven, San Diego, Karlsruhe and CERN. The largest implementation, the CDF system at FERMI, provides 150 TeraBytes of disk space and delivers up to 50 TeraBytes/day to its clients.

Sites using the LCG dCache distribution are more or less operating the cache as black box and little knowledge is available about customization and enhanced features.

This presentation is therefor intended to make non dCache users curious and enable dCache users to better integrate dCache into their site specific environment. Beside many other topics, paper will touch on the possibility of dCache to closely cooperate with tertiary storage systems, like Enstore, Tsm and HPSS. It will describe the way dCache can be configured to attach different pool nodes to different user groups but let them all use the same set of fall back pools. We will explain how dCache takes care of dataset replication, either by configuration or by automatic detection of data access hot spots. Finally we will report on ongoing development plans.