

Integration of the Trigger and Data Acquisition Systems in ATLAS

M Abolins⁴², P Adragna⁵⁷, E Aleksandrov¹⁷, I Aleksandrov¹⁷,
A Amorim³⁶, K Anderson¹³, X Anduaga³⁷, I Aracena⁶², L Asquith⁶⁹,
G Avolio¹¹, S Backlund¹¹, E Badescu⁹, J Baines⁵⁸, P Barria^{61, 29},
R Bartoldus⁶², S Batreanu^{9, 11}, H P Beck⁶, C Bee⁴⁶, P Bell⁴⁴,
W H Bell²¹, M Bellomo⁵⁶, K Benslama⁵⁹, D Berge¹¹, N Berger^{34, 27},
T Berry²⁵, M Biglietti⁵², R E Blair¹, A Bogaerts¹¹, T Bold⁶⁸,
M Bosman⁵, J Boyd¹¹, B Breliev⁴⁸, D Burckhart-Chromek¹¹,
C Buttar²¹, M Campanelli⁴², M Caprini⁹, G Carlino⁵², D Casadei⁵¹,
P Casado⁵, G Cataldi³⁹, D Cimino²⁸, M Ciobotaru^{9, 68, 11},
D Clements²¹, A Cocco²⁰, Conde Muino³⁶, F Conventi⁵²,
A Corso-Radu⁶⁸, M J Costa⁷⁰, R Coura Torres⁶⁰, R Cranfield⁶⁹,
K Cranmer⁵¹, G Crone⁶⁹, M Dam⁵⁰, D Damazio⁴, I Dawson⁶⁵,
J Dawson¹, J De Almeida Simoes⁴⁰, S De Cecco^{61, 29}, A De Santo²⁵,
M DellaPietra⁵², P-A Delsart⁴⁸, S Demers⁶², B Demirkoz¹¹,
A Di Mattia⁴², C Dionisi^{61, 29}, P, R Djilkibaev⁵¹, R Dobinson¹¹,
M Dobson¹¹, A Dotti²⁸, M Dova³⁷, G Drake¹, M-A Dufour⁴⁷,
S Eckweiler⁴³, W Ehrenfeld¹⁵, T Eifert¹⁹, N Ellis¹¹, D Emelianov⁵⁸,
D Enoque Ferreira de Lima⁶⁰, Y Ermoline⁴², I Eschrich⁶⁸, K Facius⁵⁰,
S Falciano²⁹, P Farthouat¹¹, E Feng¹³, J Ferland⁴⁸, R Ferrari⁵⁶,
M L Ferrer¹⁸, G Fischer²⁶, T Fonseca-Martin¹¹, D Francis¹¹,
S Gadomski⁶, H Garitaonandia Elejabarrieta⁵, G Gaudio⁵⁶,
O Gaumer¹⁹, S George²⁵, S Giagu^{61, 29}, R Goncalo²⁵, B Gorini¹¹,
E Gorini^{64, 39}, S Gowdy⁶², I Grabowska-Bold¹¹, S Grancagnolo^{64, 39},
B Green²⁵, S Haas¹¹, W Haberichter¹, H Hadavand⁶³, C Haeberli⁶,
J Haller^{23, 15}, A Hamilton¹⁹, J R Hansen⁵⁰, M Hauschild¹¹,
R Hauser⁴², S Head⁴⁴, S Hillier⁷, A Hoecker¹¹, T Hryn'ova¹¹,
R Hughes-Jones⁴⁴, J Huston⁴², J Idarraga⁴⁸, O Igonkina⁵⁴, M Inada³³,
V Jain³⁰, K Johns³, M Joos¹¹, S Kama¹⁵, N Kanaya³³, A Kazarov⁵⁵,
R Kehoe⁶³, G Khorauli⁵⁹, G Kieft⁵³, G Kilvington²⁵, J Kirk⁵⁸,
H Kiyamura³³, S Klous⁵³, S Kolos⁶⁸, T Kono¹¹, N Konstantinidis⁶⁹,
K Korcyl¹⁴, K Kordas⁶, V Kotov¹⁷, A Krasznahorkay^{11, 16},
T Kubota⁶⁷, A Kugel⁴⁵, D Kuhn³¹, H Kurasige³³, T Kuwabara⁶⁷,
R Kwee⁷⁴, A Lankford⁶⁸, T LeCompte¹, L Leahu^{9, 11}, M Leahu⁹,
F Ledroit²², G Lehmann Miotto¹¹, X Lei³, D Lellouch⁷², M Leyton³⁵,
S Li¹⁵, H Lim¹, T Lohse²⁶, M Losada⁸, C Luci^{61, 29}, L Luminari²⁹,
L Mapelli¹¹, B Martin¹¹, B T Martin⁴², F Marzano²⁹, J Masik⁴⁴,
T McMahon²⁵, R Mcpherson⁷¹, M Medinnis¹⁵, C Meessen⁴⁶,
C Meirosu⁹, A Messina¹¹, A Mincer⁵¹, M Mineev¹⁷, A Misiejuk²⁵,
K Moenig⁷⁴, F Monticelli³⁷, A Moraes⁴, D Moreno⁸, P Morettini²⁰,

R Murillo Garcia¹¹, K Nagano³², Y Nagasaka²⁴, A Negri⁶⁸,
P Nemethy⁵¹, A Neusiedl⁴³, A Nisati^{61, 29}, M Nozicka¹⁵, C Omachi³³,
B Osculati²⁰, C Osuna⁵, C Padilla¹¹, N Panikashvili⁶⁶, F Parodi²⁰,
E Pasqualucci^{61, 29}, T Pauly¹¹, V Perera⁵⁸, E Perez⁵, V Perez Reale¹¹,
J Petersen¹¹, R Piegai¹⁰, J Pilcher¹³, G Pinzon⁸, B Pope⁴²,
C Potter⁴⁷, M Primavera³⁹, V Radescu¹⁵, S Rajagopalan⁴,
P Renkel⁶³, M Rescigno⁶¹, S Rieke⁴³, C Risler²⁶, I Riu⁵,
S Robertson⁴⁷, C Roda²⁸, D Rodriguez⁸, Y Rogriquez⁸, Y Ryabov⁵⁵,
P Ryan⁴², D Salvatore¹², C Santamarina⁴⁷, C Santamarina Rios⁴⁷,
D Scannicchio⁵⁶, D A Scannicchio⁵⁶, C Schiavi²⁰, J Schlereth¹,
I Scholtes¹¹, D Schooltz⁴², W Scott⁵⁸, E Segura⁵, N Shimbo³³,
A Sidoti²⁹, G Siragusa^{64, 39}, S Sivoklokov⁴⁹, J E Sloper¹¹,
M Smizanska³⁸, I Soloviev¹¹, R Soluk², S Spagnolo^{64, 39}, R Spiwoks¹¹,
S Stancu^{9, 68, 11}, P Steinberg⁴, J Stelzer¹¹, A Stradling⁷³, D Strom⁵⁴,
J Strong²⁵, D Su⁶², S Sushkov⁵, M Sutton⁶⁹, T Szymocha¹⁴,
S Tapprogge⁴³, S Tarem⁶⁶, Z Tarem⁶⁶, P Teixeira-Dias²⁵,
K Tokoshuku³², E Torrence⁵⁴, F Touchard⁴⁶, L Tremblet¹¹,
M Tripiana³⁷, G Usai¹³, B Vachon⁴⁷, W Vandelli¹¹, A Ventura³⁹,
V Vercesi⁵⁶, J Vermeulen⁵³, J Von Der Schmitt⁴¹, M Wang⁵⁹,
A Watson⁷, T Wengler⁴⁴, P Werner¹¹, S Wheeler-Ellis⁶⁸, F Wickens⁵⁸,
W Wiedenmann⁷³, M Wielers⁵⁸, H Wilkens¹¹, F Winklmeier¹¹,
E-E Woehrling⁷, S-L Wu⁷³, X Wu¹⁹, S Xella⁵⁰, Y Yamazaki³²,
M Yu⁴⁵, F Zema¹¹, J Zhang¹, L Zhao⁵¹, H Zobernig⁷³, A dos Anjos⁷³,
M zur Nedden²⁶, E Özcan⁶⁹ and G Ünel^{68,11}

E-mail: benedetto.gorini@cern.ch

¹ Argonne National Laboratory, Argonne, Illinois

² University of Alberta, Edmonton

³ University of Arizona, Tucson, Arizona

⁴ Brookhaven National Laboratory (BNL), Upton, New York

⁵ Institut de Física d'Altes Energies (IFAE), Universitat Autònoma de Barcelona, Bellaterra (Barcelona)

⁶ Laboratory for High Energy Physics, University of Bern, Bern

⁷ School of Physics and Astronomy, The University of Birmingham, Birmingham

⁸ Universidad Antonio Narino, Bogotá, Colómbia

⁹ National Institute for Physics and Nuclear Engineering, Institute of Atomic Physics, Bucharest

¹⁰ University of Buenos Aires, Buenos Aires

¹¹ European Laboratory for Particle Physics (CERN), Geneva

¹² Dipartimento di Fisica dell' Università della Calabria e I.N.F.N., Cosenza

¹³ University of Chicago, Enrico Fermi Institute, Chicago, Illinois

¹⁴ Institute of Nuclear Physics, Polish Academy of Sciences, Cracow

¹⁵ Deutsches Elektronen-Synchrotron (DESY), Hamburg

¹⁶ University of Debrecen

¹⁷ Joint Institute for Nuclear Research, Dubna

¹⁸ Laboratori Nazionali di Frascati dell' I.N.F.N., Frascati

¹⁹ Section de Physique, Université de Genève, Geneva

²⁰ Dipartimento di Fisica dell' Università di Genova e I.N.F.N., Genova

²¹ Department of Physics and Astronomy, University of Glasgow, Glasgow

²² Laboratoire de Physique Subatomique et de Cosmologie de Grenoble (LPSC),

IN2P3-CNRS-Université Joseph Fourier, Grenoble

²³ University of Hamburg, Germany

²⁴ Hiroshima Institute of Technology, Hiroshima

²⁵ Department of Physics, Royal Holloway and Bedford New College, Egham

²⁶ Institut für Physik, Humboldt Universität, Berlin

²⁷ Institut National de Physique Nucleaire et de Physique des Particules
²⁸ Dipartimento di Fisica dell' Università di Pisa e I.N.F.N., Pisa
²⁹ I.N.F.N. Roma
³⁰ Indiana University, Bloomington, Indiana
³¹ Institut für Experimentalphysik der Leopold-Franzens-Universität Innsbruck, Innsbruck
³² KEK, High Energy Accelerator Research Organisation, Tsukuba
³³ Kobe University, Kobe
³⁴ Laboratoire d'Annecy-le-Vieux de Physique des Particules (LAPP), IN2P3-CNRS, Annecy-le-Vieux
³⁵ Lawrence Berkeley Laboratory and University of California, Berkeley, California
³⁶ Laboratorio de Instrumentacao e Fisica Experimental, Lisboa
³⁷ National University of La Plata, La Plata
³⁸ Department of Physics, Lancaster University, Lancaster
³⁹ Dipartimento di Fisica dell' Università di Lecce e I.N.F.N., Lecce
⁴⁰ University Católica-Figueira da Foz and University Nova de Lisboa, Lisbon
⁴¹ Max-Planck-Institut für Physik, München
⁴² Michigan State University, Department of Physics and Astronomy, East Lansing, Michigan
⁴³ Institut für Physik, Universität Mainz, Mainz
⁴⁴ Department of Physics and Astronomy, University of Manchester, Manchester
⁴⁵ Lehrstuhl für Informatik V, Universität Mannheim, Mannheim
⁴⁶ Centre de Physique des Particules de Marseille, IN2P3-CNRS, Marseille
⁴⁷ Department of Physics, McGill University, Montreal
⁴⁸ University of Montreal, Montreal
⁴⁹ Moscow State University, Moscow
⁵⁰ Niels Bohr Institute, University of Copenhagen, Copenhagen
⁵¹ Department of Physics, New York University, New York
⁵² Dipartimento di Scienze Fisiche, Università di Napoli 'Federico II' e I.N.F.N., Napoli
⁵³ FOM - Institute SAF NIKHEF and University of Amsterdam/NIKHEF, Amsterdam
⁵⁴ University of Oregon, Eugene, Oregon
⁵⁵ Petersburg Nuclear Physics Institute (PNPI), St. Petersburg
⁵⁶ Dipartimento di Fisica Nucleare e Teorica dell' Università di Pavia e I.N.F.N., Pavia
⁵⁷ Department of Physics, Queen Mary and Westfield College, University of London, London
⁵⁸ Rutherford Appleton Laboratory, Chilton, Didcot
⁵⁹ University of Regina, Regina
⁶⁰ Universidade Federal do Rio de Janeiro, COPPE/EE/IF, Rio de Janeiro
⁶¹ Dipartimento di Fisica dell' Università di Roma I 'La Sapienza'
⁶² Stanford Linear Accelerator Center (SLAC), Stanford
⁶³ Department of Physics, Southern Methodist University, Dallas, Texas
⁶⁴ Università degli Studi del Salento
⁶⁵ Department of Physics, University of Sheffield, Sheffield
⁶⁶ Department of Physics, Technion, Haifa
⁶⁷ International Center for Elementary Particle Physics, University of Tokyo, Tokyo
⁶⁸ University of California, Irvine, California
⁶⁹ Department of Physics and Astronomy, University College London, London
⁷⁰ Instituto de Fisica Corpuscular (IFIC) Universidad de Valencia
⁷¹ University of Victoria, Victoria
⁷² Department of Particle Physics, The Weizmann Institute of Science, Rehovot
⁷³ Department of Physics, University of Wisconsin, Madison, Wisconsin
⁷⁴ Deutsches Elektronen-Synchrotron (DESY), Zeuthen

E-mail: benedetto.gorini@cern.ch

Abstract. During 2006 and the first half of 2007, the installation, integration and commissioning of trigger and data acquisition (TDAQ) equipment in the ATLAS experimental area have progressed. There have been a series of technical runs using the final components of the system already installed in the experimental area. Various tests have been run including ones where level 1 preselected simulated proton-proton events have been processed in a loop mode through the trigger and dataflow chains. The system included the readout buffers containing the events, event building, level 2 and event filter trigger algorithms. The scalability of the system

with respect to the number of event building nodes used has been studied and quantities critical for the final system, such as trigger rates and event processing times, have been measured using different trigger algorithms as well as different TDAQ components. This paper presents the TDAQ architecture, the current status of the installation and commissioning and highlights the main test results that validate the system.

1. Introduction

The ATLAS experiment[1] is one of the four experiments aimed at studying high-energy particle interactions at the Large Hadron Collider (LHC), that is under construction at CERN in Geneva and is scheduled to start to operate in the year 2008. At present the installation of the different components of the ATLAS detector is being completed in the underground cavern and the commissioning process has started.

The ATLAS TDAQ has been designed to take maximum advantage of the physics nature of very high-energy hadron interactions. In particular, the Region-of-Interest (RoI) mechanism is used to minimise the amount of data needed to calculate the trigger decisions thus reducing the overall network data traffic considerably.

The selection and data acquisition software has been designed in-house, based on industrial technologies (such as CORBA[2], CLIPS[3] and Oracle). Software releases are produced on a regular basis and exploited on a number of test beds as well as for detector data taking in test labs and test beams.

The final system will consist of a few thousands processors, interconnected by a complex Gbit Ethernet network with multiple switches.

2. System architecture

The ATLAS TDAQ is based on three levels of online event selection. Figure 1 shows the principal components of the system and the expected event rate at each stage.

The first level trigger (L1)[4] provides an initial reduction of a factor $\sim 10^3$ of the event rate starting from the 40 MHz nominal bunch crossing rate of the LHC, based on information from the muon trigger chambers and on reduced-granularity information from the calorimeters.

During the latency of the L1 trigger selection algorithms (up to $2.5 \mu s$), the complete event data is kept in the pipeline memories of the detector front-end electronics. Only the data for the events selected by the L1 trigger is then transferred from these front-end memories into the readout buffers (ROBs) contained in the readout system units (ROs), where it is temporarily stored and provided on request to the following stages of event selection. The data from the large number of detector readout channels is multiplexed into 1600 data fragments by the detector-specific readout drivers (RODs) and each of these fragments is sent for storage to an individual ROB.

The maximum rate of events accepted by the L1 trigger that can be handled by the ATLAS front-end systems is limited to 75kHz, but an upgrade to 100kHz is considered for a later phase. Trigger studies estimates the L1 rate required to meet the ATLAS physics program needs, to be about a factor two lower than this limit.

For every accepted event, the L1 system produces the “Region of Interest” (RoI) information, which includes the positions of all the identified interesting objects in units of pseudo-rapidity (η) and azimuthal angle (ϕ). This information is sent by the different elements of the L1 trigger system to the RoI builder (ROIB), which assembles it into a unique data fragment and sends it to a Level 2 supervisor (L2SV). The L2SVs are the steering elements of the second trigger level (L2), which is designed to provide an additional factor 20-30 in reduction power. They receive

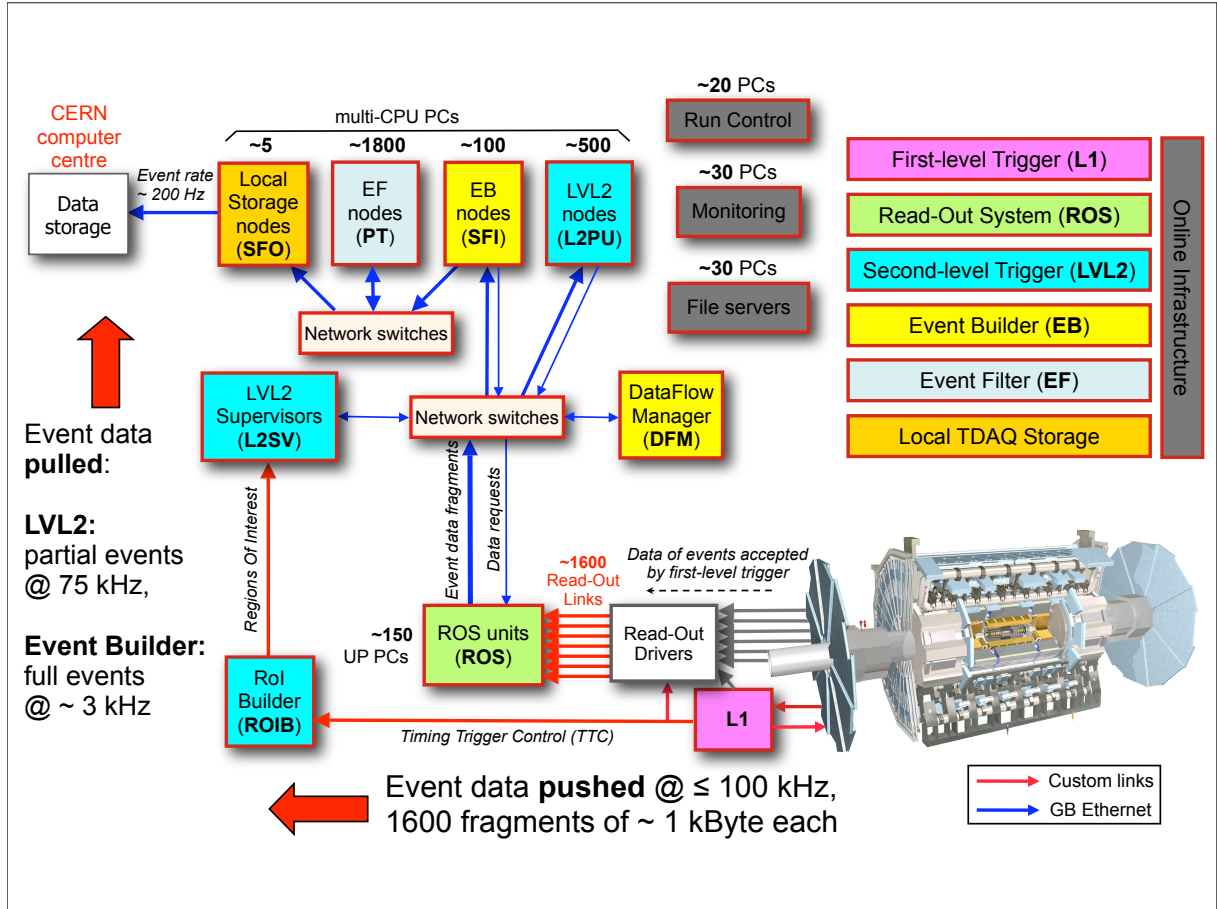


Figure 1. Diagram of principal components of the ATLAS TDAQ system.

the RoI information, assign the events to one of the processing units (L2PUs) running on a L2 node, and handle the results of the selection algorithms.

To provide the requested reduction power the L2PUs need to access detailed information from all the ATLAS detector elements (muon system, calorimeters and inner detector). To minimise the data transfers required at this early stage, the L2PUs retrieve only the few data fragments related to the geographical addresses of the interesting objects identified by the L1 (1-2 % of the total data volume). To do so it uses the RoI information received by the L2SV to identify and access only the few ROBs containing the relevant data fragments. A fast identification of the relevant ROBs is made possible by the fact that there is simple and fixed correspondence between the RoI regions and the ROBs, as each of them always receive data fragments from the same specific detector front-end modules. The L2 system is really the most characteristic element of the ATLAS architecture, and provides detailed selection power before the full event-building and consequently reduces the overall dataflow power needs.

The results of the L2 algorithms are sent by the L2SVs to the dataflow manager (DFM), which assigns the accepted events to the event building nodes (SFIs) according to load-balancing criteria. The SFIs collect the data fragments related to any assigned event from all the ROBs and assemble them in a unique event fragment. The expected rate of events at this stage is 3.5 kHz, that given a mean ATLAS event size of 1.6 Mbyte, corresponds to a total throughput of about 6 GByte/s out of the event building system.

The resulting complete event fragments are then sent to the event filter nodes, where they are assigned to a processing task (PT) running on an Event Filter Processor (EFP) for the last selection stage. The accepted events are finally sent to the output nodes (SFOs) to be temporarily buffered on local disks and transferred to the CERN computing center for permanent recording on mass storage. At this stage the rate of events is expected to be 0.2 kHz, i.e. more than a factor 10^5 lower than the original LHC bunch-crossing rate.

The L2 and EF filtering stages are collectively referred to as high level trigger (HLT).

The DFM also manages the list of events that can be removed from the dataflow system, as they have either been rejected by the L2 or received by an EFP, and periodically sends to the ROB the list of data fragments to be released.

3. System implementation

The ROB is implemented into custom PCI cards (ROBINS)[5] each housing 3 independent buffers. The ROBINS are themselves installed in PCs each one corresponding to a ROS. The connection between the RODs (detector specific) and the ROB is implemented with point-to-point optical readout links (ROLs) conforming to the S-LINK specification [3] and providing individual data throughput of up to 160 MByte/s. A ROS typically houses 4 ROBINS, for a total of 12 ROB, and handles the data requests (from L2PUs and SFIs) for all of them through its network interfaces. Upon reception of a data request, the ROS application collects the relevant data fragments from the ROBIN modules through the PCI bus (from few selected ROB for L2PU requests or from all of them for SFI requests), combines them into a unique ROS data fragment and sends it back to the requester. The total number of ROSs and ROBINS for the current ATLAS detector architecture is respectively 149 and 549 but few other units may be needed to handle the luminosity detectors that are being developed.

The RoIB is implemented as a custom VMEbus system, receiving the individual RoI information fragments and sending the combined result to the L2SVs with the same point-to-point link technology as the one used for the ROLs. The performance of communication protocol between the RoIB and the L2SVs has been measured to scale linearly with the number of L2SVs. The system can thus grow to handle different L1 rates. Three L2SVs will be required to handle the ATLAS maximum L1 rate of 75 kHz.

All other TDAQ nodes are implemented as Linux PCs running multi-threaded C++ applications. The various nodes are interconnected by multi-layer Gbit Ethernet networks and a custom message passing protocol has been developed to manage the data movements. The size of the final system will be largely dominated by the number of HLT processing nodes. In the current baseline implementation there will be ~ 500 L2 nodes and ~ 1800 EF nodes. Each of these HLT nodes will be implemented as a quad-core dual CPU machine with ~ 2 GHz of clock speed, running eight independent filtering process each (i.e. one per core). Considering that each process filters one event at the time one can derive that for a L1 rate of 75 kHz the average processing time per event shall not exceed 40 ms for the L2 and 4 s for the EF. These numbers are different from the ones originally presented in the TDAQ Technical Design Report (TDR)[6] as at the time when it was written we foresaw, by a straight extrapolation of Moore's CPU scaling law, a baseline HLT node with 2 CPUs and clock speed of 8 GHz each. Another significant difference compared to the architecture deployment presented in the TDR is that we are going to connect 900 HLT nodes to both the L2 and EF networks, providing a way to dynamically configure the relative allocation of L2 and EF processing power.

A complex online software infrastructure has been developed to configure, control, and monitor such a large system. An estimated number of 80 linux PCs is allocated for running the infrastructure.

Coherent software releases containing both the dataflow applications and the infrastructure components, are produced several times per year.

4. Status of installation and commissioning

A significant fraction of the system described in the previous sections has been installed and commissioned as of today. In particular all of the Readout system and a large fraction of the Event Building system are installed, commissioned and regularly operated either for TDAQ tests or to support detector commissioning and calibration operations. The completion of the procurement and installation of the HLT processing nodes is scheduled to match the LHC schedule and the plans for ATLAS data taking.

4.1. Readout system

The total number of 149 ROS units foreseen for the final ATLAS setup have been installed and commissioned. All the required ~ 550 ROBINS are installed in the ROSs and most of those units are connected to the detector RODs and routinely used in commissioning runs.

Four additional ROS units, fully equipped with ROBIN modules, have been installed in the ATLAS experimental area to be used as “hot” spares. These units have been commissioned like the other ones and are kept in running conditions as well as used to generate fake data for TDAQ tests.

Most of the ROS units have been intensively used during the last year as “all-in-one” DAQ systems for detector commissioning, forwarding all the data received in the ROBIs to either local or remote mass storage facilities, and have proven to work very stably.

4.2. Event building

As of today 32 event building nodes (SFI) have been installed in the ATLAS experimental area and fully commissioned. This represents 30% of the foreseen final system and 60% of the system required for the initial running period in 2008.

The installed SFIs are based on a dual Opteron 252 CPU, with a clock speed of 2.6 GHz and 2 GB of RAM. Each SFI has three Gbit Ethernet network interfaces, two on-board and one on a PCI-express NIC, that are dedicated to control, data input and data output.

In addition to the SFI nodes, twelve manager nodes (DFM) have been installed. The DFMs are based on the same kind of PC as the SFIs. As described in section 2 the final system will only need one manager application. More have been installed to support system partitioning, i.e. the possibility of taking data independently and in parallel from different subsets of the ATLAS detector, either for calibration or for testing purposes.

As described in section 2 the data selected by the HLT needs to be stored locally, prior to be transferred to mass storage in the CERN computing center. This local storage functionality is provided by the SFO nodes. All the required 5 SFO nodes plus a spare have been installed and commissioned. Each node has 24 disks of 500 GB each, on three SATA RAID controllers, for a total of 12 TB per node. Considering the expected rate of accepted event for the final ATLAS and the redundancy required by the RAID protocol, the 5 SFO nodes will be able to store locally about 24 hours of data, to cope with possible interruptions of the communication between ATLAS experimental area and the CERN main site.

4.3. HLT nodes

At the time of the conference 130 HLT nodes have been installed and commissioned that represent slightly less than 6% of the final system. These nodes have been connected to both the L2 and EF networks, such that they could be used as both L2PUs or EFPs. The current installation plan foresees to upgrade the system by a factor eight by the time of the first ATLAS run, which is foreseen for the second half of 2008.

The current choice for the HLT nodes is a PC with a dual Intel E5320 CPUs. Each of the CPUs have 4 cores with a clock speed of 1.86 GHz and 1GB of memory per core. Future nodes may have different characteristics, according to the time of purchase. The idea is to purchase

HLT computing power as late as possible, following the development of the ATLAS physics needs. This approach ensures an optimal utilization of the resources, as we will always be buying state-of-the-art machines at the time we need them.

4.4. Infrastructure

To guarantee the best possible system uniformity, still following the evolution of the TDAQ software, we decided that all ROS, EB and HLT PCs as well as all the single board computers controlling the detector front-end electronics would be net-booted. At the moment 25 file servers have been deployed for the purpose, but the number is subject to increase following the global system evolution.

An additional 20 PCs have been installed to implement the control system tree, that dispatches commands and configuration data to the different nodes of the system, ensuring the global synchronization through an expert system implementing a finite state machine.

Another 32 PCs have been deployed, though only 15 have been commissioned and used at the moment, to provide monitoring functionality for all the elements of the TDAQ system as well as all the detector electronics.

Both the number of control and monitoring machines is considered to be the final one.

5. Results from system tests

Both the global architecture and the specific choice of hardware components have been tested in a number of ways in the last few years. As the results of the large number of tests performed on individual components or on small scale testbeds have already been described in detail in several conference proceedings, the present paper just provides few highlights on the tests that are being performed on the final system.

Some of the main results of performance tests performed on ad-hoc testbeds on the TDAQ system and in particular on the ROSs are reported in [7].

5.1. Event building tests

After all the successful validation tests performed in the past on the event building architecture, the major open question was to verify the scalability of the system's performance to the number of ROS and SFI nodes of the final ATLAS.

Figure 2 shows the aggregate event building throughput as a function of the number of SFI nodes, and using 124 out of the 149 final ROSs to internally generate data fragments with a total event size similar to what is expected to be produced by the ATLAS detector in final running conditions. To push the scalability study beyond the number of installed SFI machines (32) we configured few of the HLT PCs to run as SFIs. In this setup the SFI nodes were configured to dump the events after having built them, as no connection to the EF system had been configured.

The result not only shows perfectly linear scalability, but also indicates that each SFI is able to absorb as much as 114 MB/s of incoming data that is very close to the GB Ethernet theoretical limit. One can thus deduce that the ad-hoc event building protocol developed for the ATLAS TDAQ has a negligible overhead.

Tests have also been performed to study the performance drop to be expected when enabling the output from the SFI nodes to the EF system. A setup with a ratio between SFIs and connected EFP nodes close to what is foreseen for the final system, have shown a performance degradation of just $\sim 10\%$.

5.2. HLT tests

The critical parameters to measure to validate the HLT architecture and implementation, are the average latencies for event selection in the L2 and EF.




Figure 2. Scalability of the event building throughput as a function of the number of SFI nodes.

A test setup have been put together with 124 ROSs, 32 SFIs and all the 130 HLT nodes used as L2PUs. The system has been operated pre-loading on the ROS and ROIB memories the data of $\sim 4k$ simulated collisions. The Montecarlo simulation used to produced the data included a detailed description of the complete ATLAS detector.

Figure 3 shows the average event processing time for the events rejected by the L2, that constitute the majority of events given the high selection power. The result includes both the actual processing time of the algorithms and the time required to collect the ROI data from the relevant ROB.

As one can see the measured processing time is well below the required limit of 40 ms. The result has of course a large uncertainty as it is based on simulated data, but on the other hand one has to stress that all measurements were performed on the few processing nodes available today that are likely to have significantly less computing power than the large majority that will be purchased at future times. We can thus confidently conclude that the measurements indicate a good compatibility between the ATLAS requirements and the performance provided by the TDAQ system.

More detailed results on specific algorithms are reported in [8].

Similar measurements have been performed configuring the HLT nodes as EFPs (and bypassing the L2). Initial results also show a vary good compatibility between requirements and achieved performance.

5.3. Commissioning with cosmic rays data

The installed TDAQ system has also been operated to collect real detector data, triggering on cosmic-ray particles. Detailed results of these commissioning runs have been presented in this




Figure 3. Total event processing time of the L2 trigger for rejected events. The time includes both the actual algorithm processing time and the time to retrieve the ROI data from the ROBs.

conference and are reported in [9].

6. Conclusions

A large fraction of the ATLAS Trigger and DAQ system has been installed and commissioned.

The tests performed on the system indicate that both the architecture and the specific H/W deployed, meet ATLAS requirements.

The system has been successfully used for data taking in detector commissioning runs.

HLT selection algorithms have been tested with simulated collision and real cosmic rays data.

Measurements of the HLT algorithms times indicate that they are compatible with the L2 and EF time budgets.

References

- [1] General information about ATLAS can be found on <http://atlas.web.cern.ch/Atlas/index.html>.
- [2] Common Object Request Broker Architecture, see definition on <http://www.omg.org/gettingstarted/corbaFAQ.htm>.
- [3] C Language Integrated Production System, see definition on <http://clipsrules.sourceforge.net/WhatIsCLIPS.html>.
- [4] The ATLAS Collaboration, “First-level Trigger Technical Design Report”, CERN/LHCC/98-14, 1998.
- [5] B. Green et al., “ATLAS Trigger/DAQ RobIn Prototype”, IEEE Trans. Nucl. Sci., Vol. 51, 2004, pp. 465-469.
- [6] The ATLAS Collaboration, “High-level Trigger, data Acquisition and Controls, Technical Design Report”, CERN/LHCC/2003-022, 2003.
- [7] J Vermeulen et al., “ATLAS dataflow: The read-out subsystem, results from trigger and data-acquisition system testbed studies and from modeling.”, IEEE Trans.Nucl.Sci.53:912-917,2006.
- [8] T Fonseca-Martin et al., “Event reconstruction algorithms for the ATLAS trigger”, these proceedings.
- [9] J Boyd et al., “The ATLAS Trigger: Commissioning with cosmic-rays”, these proceedings.