ATLAS DETECTOR DESCRIPTION DATABASE ARCHITECTURE

A. Vaniachine#, D. Malon, ANL, Argonne, IL 60439, USA
P. Nevski, BNL, Upton, NY 11973, USA
K. Bernardet, CPPM, F-13288, Marseille, France
S. Baranov, Freiburg University, 79085 Freiburg, Germany

Abstract

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 common Interval-of-Validity database infrastructure, measures for tighter primary numbers database input control, experience with ATLAS Combined Test Beam geometry and conditions payload storage using NOVA technologies integrated with the LCG 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.

INTRODUCTION

The Large Hadron Collider (LHC) at the CERN Laboratory will become the largest scientific instrument in the world when it starts operations in 2007. An unprecedented size and scale of the large general-purpose ATLAS detector are presenting new challenges in the field of computing. In particular, tens of thousands of unique parameters – primary numbers – necessary for ATLAS detector description are clearly beyond the scope of traditional manual parameters management.

COMPONENTS OF SUCCESS

To overcome that scalability limit we developed the powerful and intelligent user services – NOVA database services – facilitating handling of numerous parameters that steer data processing and simulations. NOVA database provides a common source of primary numbers – parameters for the detector description in all ATLAS software subsystems.

The detector description knowledge grows over time across all ATLAS subsystems (Figure 1). At the time of this writing NOVA database collects 30K of unique primary numbers – doubling the amount that was reported at the CHEP'03 [1]. In 2004 NOVA database services were deployed in ATLAS to support Data Challenge 2 (DC2) and Combined Test Beam operations.

Benefits of NOVA technologies, their conceptual decomposition into services and interoperability of NOVA components are described below.

Relational Technology

NOVA database architecture is described in [1]. It is based on a scalable relational database approach from the previous prototyping project [2].

A major architectural design principle behind the detector description scalability – a multi-level semantic and structural hierarchy of the data – is well supported by the relational model. In contrast to an alternative hierarchical technology – XML, the relational database technology is capable not only to resolve traditional many-to-one relationships between the primary numbers and their aggregation (a 'structure' in the NOVA semantic vocabulary) but also resolves reverse one-to-many relationships between the unique primary number and the

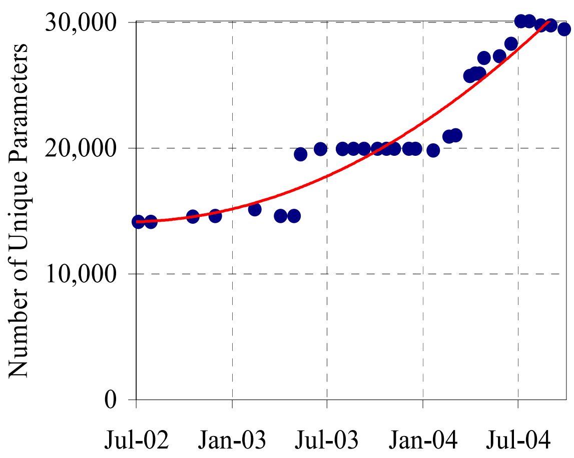

Figure 1: Growth in the primary numbers count over time manifests our deeper understanding of ATLAS detector description knowledge.

#vaniachine@anl.gov

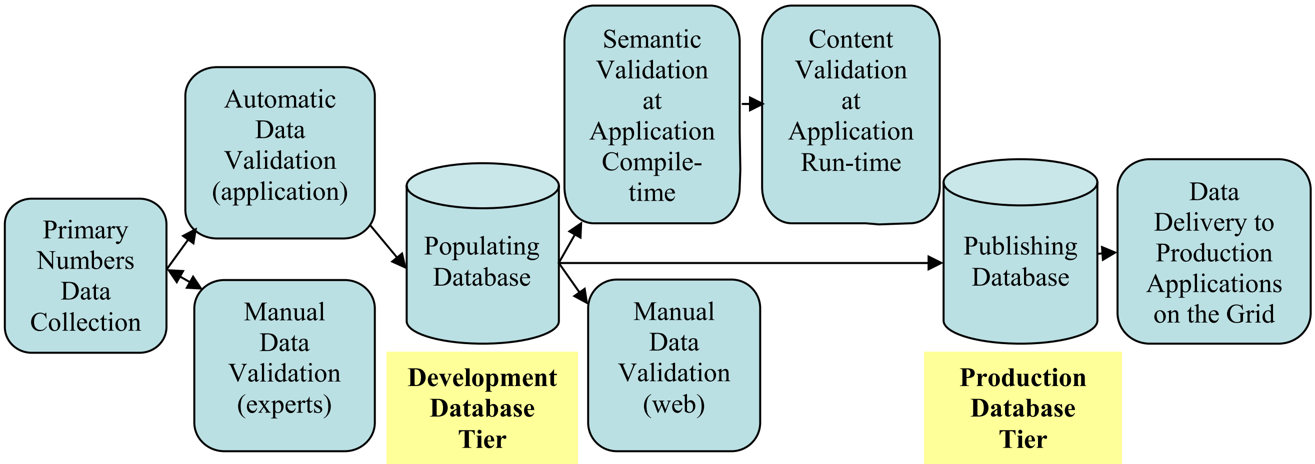

Figure 2: Workflow architecture for capturing, controlling and disseminating ATLAS detector description knowledge. Multiple validation and publishing steps are integral parts of the primary numbers data flow.

structures. The latter use case is dominant in a layered detector with repetitive geometries, where an aggregation of detector description parameters for a particular layer or element shares the same set of primary numbers.

Iterative Knowledge Capture

We have found that inherent relational database capabilities of modelling and managing relations between the primary numbers are essential to automate the laborious and iterative workflow of the detector description knowledge discovery and capture (Figure 2).

Use of core relational database technologies also enabled smooth integration of NOVA within ATLAS multi-tier database services architecture, providing controlled data transfer ('publishing') from the development server tier to the production server tier. This architecture provides measures for a tighter primary numbers database input control. Without a separation of development activities and production operations the detector description support becomes fragile.

Harnessing Complexity

For scalable management of ATLAS detector description complexity NOVA database architecture implements multiple levels of structural data hierarchy. Semantically close parameters are aggregated in arrays. A group of parameters and/or parameter arrays are further aggregated into structures. Structures of the same type can be joined in arrays. A set of structures from the same knowledge domain is collected into directories, which are further grouped in folders representing ATLAS detector subsystems (Table 1). The NOVA tag on top of the hierarchy is semantically related to a corresponding software release tag. To facilitate muon software development process multiple geometry versions are provided within the same database tag [3]. For geometry control both database tags and non-default versions are selectable via NOVA conversion service options.

Alternative technologies, e.g. XML, does not provide for relational capabilities that are instrumental for keeping

Table 1: ATLAS Detector Description Hierarchy

Entity	**Count**
Tag	79
Folder	20
Version	5
Directory	73
Structure type	340
Structure instance	15966
Parameter	29913

primary numbers count down to prevent the parameter explosion resulting from multiple geometry versions. Without relational technologies each ATLAS parameter on average will be duplicated 18.3 times complicating primary numbers management.

Facilitating Collaboration

Hierarchical NOVA architecture facilitated effective solutions and collaboration of many software developers from different sub-detector domains. Semantic compartmentalization along the multi-level subsystem hierarchy enabled simultaneous collaborative efforts of many ATLAS software developers and experts during the rapid software development cycle in preparation for the Combined Test Beam and is manifested by the extent of NOVA semantic vocabulary reuse: there are 332 different parameters named *dz*, 287 named *Z*, 166 named *Rmin*, etc.

Agile Technologies

A successful implementation of another geometry in support of the ATLAS Combined Test Beam detector also proved versatility of NOVA technologies for knowledge capture and management. Practically a whole new set of knowledge describing a very different detector was captured and persistified within the NOVA database services infrastructure in a limited period of time [4].

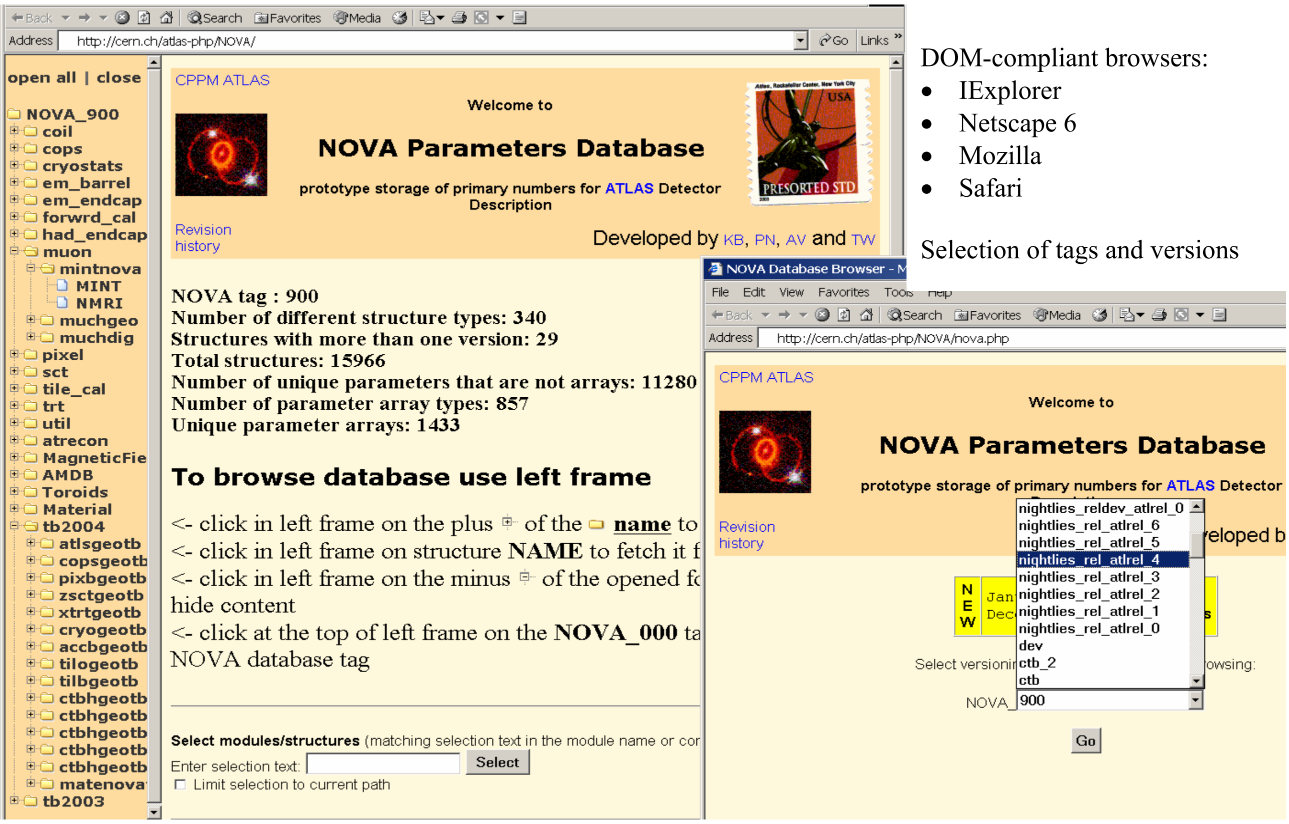

Figure 3: NOVA tools for knowledge discovery, navigation and browsing.

In addition to the Combined Test Beam primary numbers storage, the whole suite of NOVA technologies was successfully deployed in the Combined Test Beam production environment for storage and retrieval of LAr calorimeter conditions and calibration data payload via the LCG ConditionsDB interface further proving versatility and adaptability of NOVA architecture [4, 5].

Navigating Data Hierarchy

To support multiple validation steps of the primary numbers workflow new and improved user tools for knowledge discovery, navigation and browsing were deployed. Figure 3 presents screen snapshots of the new NOVA database browser for data navigation and semantic querying enhanced with NOVA tags and versions selection. A DOM-compliant browser implementation provides support for the main web browsers set: Netscape, IExplorer, Mozilla and Safari.

To assure operations of NOVA technologies for the calibrations and conditions data payload in the Combined Test Beam NOVA binary data browsing capabilities were integrated within ATLAS conditions data browser [6].

Framework Integration

A dedicated NOVA service for code generation facilitates strict C++ type-checking for all of the 30K unique parameter names at compile-time. For automatic service operation it was integrated with NICOS – ATLAS system for nightly software builds [7]. The service robustness was assured by repackaging into separate components and added locking to support simultaneous parallel release builds, builds at different sites, etc. The success of NOVA service for automatic generation of classes for persistent data access was manifested in a rapid growth in the number of NOVA object classes (Figure 4) constituting a third of ATLAS core software. To fully take advantage of Athena on-demand data retrieval architecture, a separate NOVA conversion service provides access to primary numbers. The service registers primary number objects with the transient

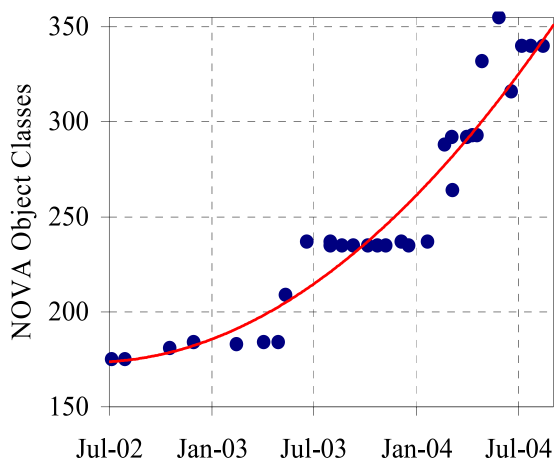

Figure 4: Growth of auto-generated NOVA object classes providing persistent-to-transient data conversion.

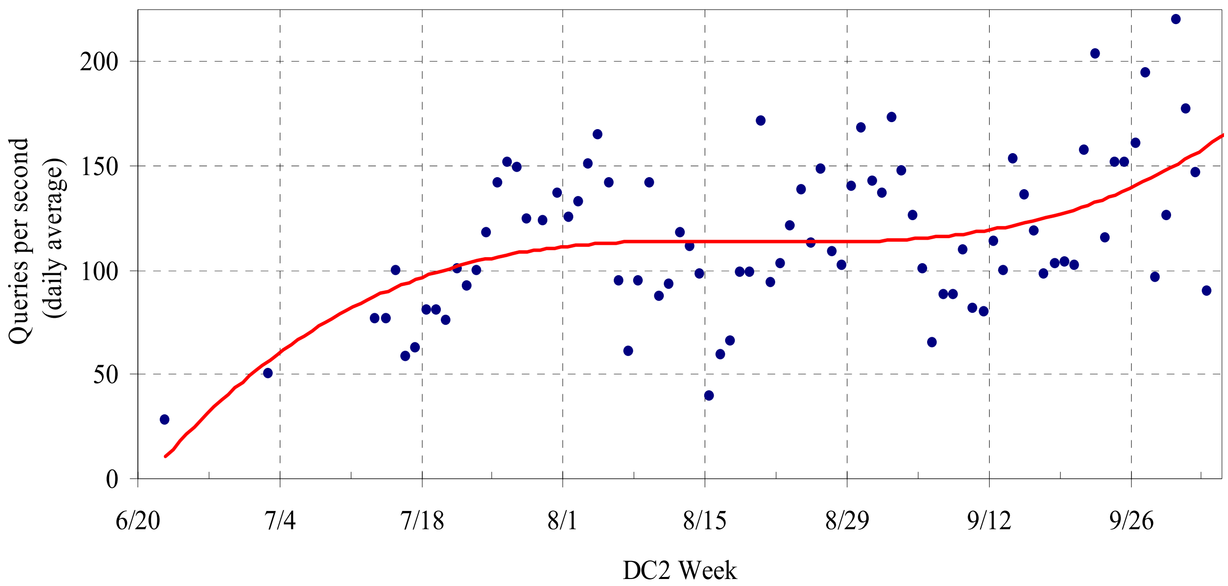

Figure 5: A growing server load from chaotic NOVA database access patterns in ATLAS DC2 production on the Grid.

detector store so that the NOVA objects can be retrieved via public StoreGate methods. NOVA conversion service is independent on database tag and NOVA objects shapes or versions. Behind the scenes, on-demand data access architecture triggers the actual database access and the appropriate conversion of persistent data to a transient object. In addition, NOVA objects retrieval is integrated with the Athena Interval-of-Validity infrastructure of the LCG ConditionsDB [1].

As a result of an increased access to NOVA services in ATLAS, limitations were found in the initial NOVA data retrieval implementation in the Athena framework where the on-demand retrieval architecture provide no well-defined place for managing input resources, like the multiple database connections. The detailed investigation and NOVA users' reports helped to isolate a potential bottleneck and provided multiple solutions: application resource pooling on a client-side, server-side timeouts, and connection management by the conversion service. All of these solutions were implemented in time for the NOVA services deployment in the production environment of the Combined Test Beam and DC2.

Validating Computing Model in DC2

Scalability of NOVA database services approach was proven through a successful deployment in the production environment of ATLAS Data Challenge 2 operations.

Most of the DC2 database-resident data flow comes from NOVA providing parameters for tens of thousands of jobs for simulation, digitization and reconstruction data transformations. This data flow is monitored to probe database services limits and identify potential bottlenecks in servicing chaotic access patterns observed in DC2 production on a federation of three computational grids: LCG, Grid3 and Nodugrid (Figure 5). In DC2 production environment the failure rate from NOVA database access on the Grid was found to be less then 0.1%.

Future Challenges

In preparation for future challenges ATLAS detector description knowledge, structural hierarchy and semantic vocabulary captured in NOVA are transitioned to new Oracle database supported by the CERN IT/DB group. Using new tools generating SQL scripts for Oracle data input, most of NOVA structures are now stored in the new Oracle database.

ACKNOWLEDGEMENTS

We thank all of our ATLAS collaborators and, in particular, developers from all software sub-systems. We also wish to thank J. Boudreau, A. Farilla, S. Goldfarb, R. Hawkings, H. Ma, D. Pomarede, L. Price, J. Proudfoot, A. Rozanov, V. Tsulaia and T. Wenaus for numerous discussions, interest and support of this work. Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under contract W-31-109-Eng-38.

REFERENCES

[1] A. Vaniachine et al., CHEP'03, La Jolla, USA, 2003, eConf C030324:MOKT006 [cs.db/0306103]

[2] P. Nevski, A. Vaniachine, T. Wenaus, CHEP'2000, Padova, Italy, 2000, Proceedings, p. 156.

[3] D. Pomarede et al., DPF'04, Riverside, USA, 2004, Proceedings, paper 372.

[4] A. Farilla et al., CHEP'04, Interlaken, Switzerland, 2004, Proceedings, paper 275.

[5] A. Amorim et al., CHEP'04, Interlaken, Switzerland, 2004, Proceedings, paper 264.

[6] A. Amorim et al., CHEP'04, Interlaken, Switzerland, 2004, Proceedings, paper 263.

[7] A. Undrus, CHEP'03, La Jolla, USA, 2003, eConf C0303241:TUJT006.