Metrics Correlation and Analysis Service (MCAS)

Andrew Baranovski [*]**, Dave Dykstra, Gabriele Garzoglio, Ted Hesselroth, Parag Mhashilkar, Tanya Levshina**

Fermi National Accelerator Laboratory, Batavia, IL, USA

{abaranov, dwd, garzoglio, tdh, parag, tlevshin}@fnal.gov

Abstract. The complexity of Grid workflow activities and their associated software stacks inevitably involves multiple organizations, ownership, and deployment domains. In this setting, important and common tasks such as the correlation and display of metrics and debugging information (fundamental ingredients of troubleshooting) are challenged by the informational entropy inherent to independently maintained and operated software components. Because such an information "pond" is disorganized, it a difficult environment for business intelligence analysis i.e. troubleshooting, incident investigation and trend spotting. The mission of the MCAS project is to deliver a software solution to help with adaptation, retrieval, correlation, and display of workflow-driven data and of type-agnostic events, generated by disjoint middleware.

1. Introduction

The Grid is a common paradigm for sharing distributed resources. With the paradigm now arguably surpassing one decade of adoption by a wide variety of communities, the amount of information provided by distributed services, such as monitoring, discovery, troubleshooting, accounting, auditing, etc. is becoming increasingly difficult to manage in a coherent fashion. Because of the disjoint nature of all these data sources, the aggregation, transformation, and display of this distributed information is particularly challenging. To address this problem, the Fermi National Accelerator Laboratory has initiated the Metrics Correlation and Analysis Service (MCAS) project [MCAS].

A core business idea of this project is to factor out presentation and business analysis logic from available monitoring solutions into a standalone model, supporting common standards in data manipulation and presentation. The MCAS project prototyped several services, which rely on common techniques for data and information display aggregation. In particular, we've chosen portlet [Portlet] technology to compose troubleshooting and metric analysis dashboard, and XAware [XAware] to drive the data-integration model.

Section 2 introduces the problem of analyzing coherently disjoint data sources and proposes data co-display and data integration as possible solutions. Section 3 discusses advantages and disadvantages of technologies for data co-display.

---

[*] To whom any correspondence should be addressed.

2. Uniform analysis of uniform data

Typically, operational problems cannot be easily visualized by the use of one particular way of displaying information. In fact, as part of the decision-making process the experts are using a variety of tools, which report data in ways that best emphasize patterns of previously addressed problems.

This variety is built on technologies that often incorporate incompatible data input and produce output that cannot be used within the same context. These incompatibilities restrict our ability to create cross-analyses between data domains of similar origin. A solution to overcoming such limitations consists in promoting intuitive ways to process and represent underlying monitoring or diagnostic information.

Phase 1 of the MCAS project has focused on a minimalist approach to define data though XML and to build a foundation to support complicated scenarios of metrics analysis and presentation. For this phase, the project has focused on creating a toolkit of display and analysis tools for presentation and co-display of data. The idea of the toolkit is to factor typical troubleshooting practices into components that can be hosted by a common container environment. The design and implementation of the delivered software is based on concepts of information display and data integration.

3. Data co-display

Data co-display is one of the simplest ways to cross-analyze data or to determine heuristic correlations between data. In the MCAS project data co-display is implemented by aggregating independently designed and managed frames inside one Web browser window. Each frame renders some aspect of the input data. This design decision is fundamental for the project. Not does it decouple development and testing of the display components, but it also allows users to choose and change the presentation layout, along with configuration details of each component, in a way that best represents the state of the system. As such, we have evaluated two technologies that implement the design idea of content aggregation: web widgets and JSR 168 Java Portlets.

Both web widgets and JSR 168 Java Portlets offer developers a framework for writing web code that can coexists in the context of a single web page. Such framework is implemented through hosting environments and, ultimately, allows developers to plug in code into service containers (iGoogle, Netvibe , MySpace). The function of the container is to augment the developer's code with features common across all elements of the display. These features include widgets setup and display position, configuration editing, help, and persistence of configuration. Web widget containers generate the front page using the developer's scripts together with the framework code required to support common functions. Finally, the generated code is rendered by the browser.

The web widget approach is extremely scalable, as it is entirely client-based and does not require a server side context. The significant drawback of this technology is in its dependence on the provider's API and, consequently, vendor lock-in for further hosting and support. This problem, however, can be mitigated by thin sub-framework wrappers, which isolate project-wide widget behavior.

In contrast to web widgets, JSR 168 Java Portlet specification is entirely server-side based. The specification sets forth a standard of development for java servlets, which generate independent, dynamic content. This context can be handled by the hosting environment for common management, configuration, and look and feel policies. Because the portlet specification is server-based, its major deficiency is scalability and the inherent complexity of the code. We believe that these weaknesses, however, are still preferable to vendor lock-in, the principal disadvantage of web widgets. In addition, the specification clearly benefits from an open standard and a variety of implementations, which are available as community driven products (JBoss[JBoss], Pluto[Pluto], eXo[eXo]). In the MCAS project we have chosen to use java portlet specifications implemented by the JBoss portal as the primary mean to compose and render project's user interfaces.

4. Data integration

Our goal is to provide unified data analysis capabilities to a disjoint, format-incompatible universe of datasets. We approach this problem through the unification of the data itself. This unification can be

accomplished via data format and semantic transformation. These transformations should make the transformed data compatible with display or analysis tools, thus enabling the analysis of the unified data. The unification in format and semantics decouples the design of the analysis software from the details of original input data. This also allows the reuse of the transformed data in contexts not originally envisioned by the project, a major benefit in enabling the utility of the solution for perhaps several years. This approach, however, has the high startup cost of developing transformations for the several data formats, while providing a usable model for describing the transformed data.

We have given the reasons why the data integration layer must be isolated and designed-in to a system which will function with diverse data sets. The question remains as to what format the unification (or data description language) must adhere to. In the MCAS project we have decided against attempting to set a specific format that fits all possible recommendations for the problem, rather we have chosen to follow a recommendation that can capture common traits without unreasonably limiting future choices.

In this context we have isolated two criteria: support for structured data and support for navigation of structured data. For the MCAS project, the first criteria is important for preserving the semantics of input . The second criteria is critical for enabling formalism in transformation of that input. Among a variety of recommendations and available third party implementations: value pair, CORBA, ad hoc binary, and XML; -XML fits the task the best. XML has naturally became the backbone supporting data unification needs as well as the only means of describing transformation workflows necessary to do that unification.

5. Architecture

The MCAS system workflow is organized between four players – user, content management system, data integration layer and data sources.

A starting point of the project is a portal or dashboard. This page has a unique address and typically offers a default view of the system which user would like to see first. The content of the page is supported by the Content Management System (CMS). The CMS is responsible for building composition of independent interface elements called "Portlets" using

1) The address of the user request
2) Static configuration parameters of the individual portlet
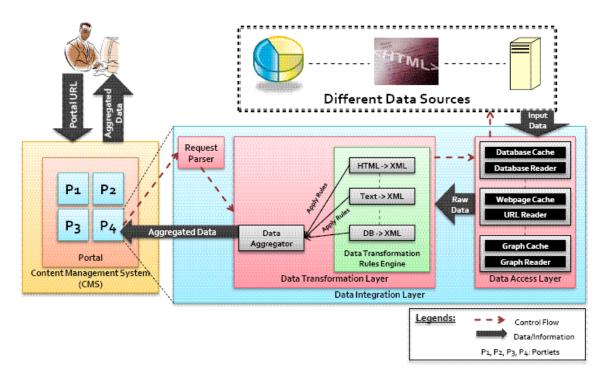3) Run time information provided by the data integration layer

Figure 1MCAS Architecture

5.1. *The JBoss Content Management Environment*

CMS renders each portlet (P1, P2, P3, and P4 as shown in the figure **1**) independently. Each portlet can be autonomously designed with some unique perspective of a particular system aspect. A collection of different portlets describing different perspective of the system, can be put together to form a dashboard. The dashboard view reflects the state of the entire system at a given time. The MCAS toolkit is based on high level data presentation routines from Google, RRD [RRD] and other lower level JavaScript codes developed in-house. The toolkit accepts data represented in standard XML schema as its input. Each portlet is rendered using information provided by the data integration layer.

The data integration layer accesses set of data sources and uses a collection of rules to transform and aggregate the retrieved content. This layer is responsible for providing content display to the CMS. The output is returned synchronously to the requesting party, which may be a portlet or another data integration rule set. Content display, is expected to be available "immediately" following user's request. Hence, the data integration layer URL must have access to the data in real-time.

In our model, the data integration layer defines all endpoints providing content for the display. These endpoints may be as simple as proxies to existing web pages or may invoke complicated rules for transforming and aggregating data retrieved from other sources. The purpose of this layer is to respond to requests for data that otherwise is unavailable or in formats not applicable to the user interface. Such applicability will typically comprise requirements to data format, data quality and data availability. Ultimately, the data integration layer will re-use existing preconfigured resources to generate the response. In doing so, it may need to compromise quantity and restrict formats of data retrieved from integrand data source endpoints in order to generate a response within an acceptable time.

A data source is a primary provider or the source of the pre-processed information. It may come in various formats and support variety of communication models. In the pilot stage of the project we focus on HTTP accessible resources.

5.2. Data integration model

The data integration model uses XAware for data source access, inter component data exchange and data transformation scheduling. XAware has extensive support for structured content based on XML and fairly rich set of features for content transformation. The product offers a well defined model supported by collection of Meta-XML [Meta-XML] documents. These documents define how the data sources should participate to transform the raw data into more useable format. This model is expressed though the following three entities – the driver, the component, and the document.

The driver is the Meta document which instructs how a data source can be accessed. It also specifies the data source address URL and other parameters required to retrieve content from the data source.

The component is a logical view of the data source. It is a Meta document which defines the XML schema of the driver output and the driver input. Its purpose is to fix the data source XML schema and encapsulate necessary transformations to ensure the compliance of the raw content.

The document is the ultimate product of the XAware workflow. The document is built from the content provided by other documents or components. This model sets the structure of the output XML as well as defining the workflow of transformations for building that structure.

Each document name is mapped into URL endpoint accessible to the user through the XAware hosting environment. When user contacts such URL, the data integration layers triggers series of transformations in strict accordance to the hierarchy of Meta object references comprising the overall document. The tree starts at the user contact point and ends with the driver. The resulting document is sent back in plain text using HTTP.

All Meta documents describing the workflow are part of the hierarchy of the resulting document. This design places limitation on flexibility of the output and options to dynamically influence the transformation sequence. In the initial phase of the project, loss of these capabilities is not critical. We have decided to extensively invest into XAware features to transform informational content for presentation by the dashboard displays. However, given product's weaknesses, we are also investigating other platforms for data integration. One such alternative is Mule ESB.

5.3. Mule ESB

We have evaluated Mule ESB as the initial prototype suite for functions of the data integration layer. Attractive features of the Mule ESB are capabilities of accessing diverse data source transports, support for message based inter component exchange, concurrency, synchronicity control and staged execution scheduling. One of the major features of Mule ESB integration platform is the ability to set up data and execution flow in a transport/interface agnostic way. In particular Mule ESB offers:

1) Codes to translate, or templatized translation of data formats
2) Options to manage synchronicity with choices ranging from fully synchronous to Stage Event Driven Architecture (SEDA [SEDA]) based solutions.
3) Codes which adapt out of the box to different transports (TCP, UDP, SMTP, FTP, jdbc, etc)

5.4. Messaging

Message exchange is a clever way to decouple contexts of two programs. Messaging is a soft pattern which does not rely on a fixed db schema or file format. Rather, its tableau is focused on data transport

and synchronicity issues. Consequently, within the Mule-message enabled infrastructure there are options for using opaque payloads; modeling of data access and aggregation is separated from specifics of the type-structure of the data sources. This is in contrast to the XAware model, in which transformation is assumed and all data exchanges have predetermined schema.

Messaging and data integration models have been used in an evaluation phase of the MCAS project to connect set of message enabled services into a workflow that refactors existing informational portals. Information from a set of D0 site efficiency sensors, accessible through http, has been assembling in a Content Management System. Figure 2 below depicts the data and execution flow of the implementation, which uses formal data transformations and RRD processing engine to do splitting, rescaling and redrawing of the D0 site efficiency data.

In our Mule workflow implementation we used a SEDA based message communication model. Mule messages communicate information and trigger component invocations. The generic Mule message consists of message envelope and enclosed opaque payload.

A necessary condition for supporting opaque content of the transport message is an interface to adapt that content to a format acceptable by the Mule endpoint components – a business logic container. This functionality is a major source of flexibility that allows developers to override Mule message end point interfaces and deploy customized transformations. Examples of such end point implementations are data base, mail server, or RRD rendering engines. For the prototype evaluation phase of the project we implemented a RRD rendering engine. The engine allows a user to formulate simple data manipulations and set the parameters of the resulting display. Inside the RRD image rendering engine we were able to convert input data in the form of XML into a collection of native RRD databases. The content of those databases is then used by native RRD to perform data manipulations using a command invoked from a template language similar to the example given below.

```
ds(D0ProductionEfficiency)
ec=eff_code; ef=eff_fini;
RRD(CDEF:ec_adj=ec,0,100,
LIMIT CDEF:ef_adj=ef,0,100,
LIMIT LINE2:ec_adj#FF0000:eff_code(x100)
LINE2:ef_adj#0000FF:eff_fini(x100))
imgsize(600,300)
```

This meta code directs the workflow to access D0ProductionEfficiency time series data source, split the data source content into two streams using values of data property ("eff_code", "eff_fini") and finally instantiates actual RRD command while referencing data factorized at the previous step.

The data transformation engine is built by setting up a model to describe message-driven interactions between Mule ESB message endpoints. This particular schema is designed to execute a template-like language and has only one data source (the production efficiency endpoint). The embedded RRD template enables transformations over split data streams. The result of the transformation is a new document with an image, which is sent to a portlet instance specifically configured to interact with this data integration model.
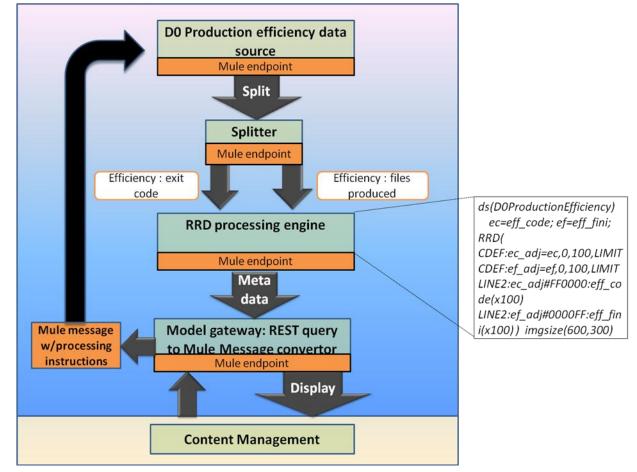
Transformation pipeline

Figure 2 An example of data integration workflow

5.5. Architecture Summary

The Data Integration layer contains definitions of all URL endpoints that provide content for the display. These endpoints may be proxies to existing web pages or they may invoke rules for transforming and aggregating data retrieved from other sources. In all cases the data results in an XML representation of the resources accessed. The output of the proxies or rules are data of the semantics and format required by the portal Content Management System, which in this case is the JBoss portlet capability. The user interface is invoked in a web browser, displaying the portlets which have been produced from the aggregated data.

Though XAware as the data integration layer is strong in transformational capabilities, the static nature of its transformation sequence has led us to evaluate Mule as the integration component. Mule's focus on transport and synchronicity issues allows separation of these factors from schema and format. We implemented a workflow in Mule to collect D0 production efficiency data, transforming it into images to be displayed via portlets.

6. Current work

At this stage of the project, the MCAS team focus is not on providing sophisticated means for analysis of distributed metrics or diagnostics data. Instead, the project is focused on development of easy to use

and intuitive tools for displaying data already reported by existing information portals. Following that path, we have already developed a substantial collection of tools that fit the monitoring needs of Fermilab's USCMS facility operations. This software enhances the capabilities of existing experiment portals by bringing the expertise we have gained from analyzing, monitoring, and troubleshooting use cases of users and operators from several different experiments. The current capabilities of the software are described in the consecutive sections.

6.1. Table view

Table view renders a data table model. Each row of the table undergoes summarization which determines the relation between columns of the model data and columns of the data displayed on the web page. The table view portlet supports custom sorting features, table size constraints, and user-assigned color coding of different weight. The example of the input data format is given below:

```
<table xmlns="http://www.fnal.gov/docs/products/mcas/TableView">
 <column name="c1">
  <row>value1</row>
 </column>
 <column name="c2">
  <row>value2</row>
 </column>
</table>
```

6.2. Bar graph

Bar graph displays a collection of indicators that visualize value pairs in relative proportions. These collections are often used to show health or performance status of the system relative to predefined metrics. This widgets idea is to help compact information collected from variety of "status" pages into one document.

6.3. Time series

Time series widgets adapt Google annotated time line code and uses its features to visualize XML documents formulated to convey data values changing in time.

6.4. Image display

The image display portlet accepts an XML document that characterizes list of URL locations for image data and generates HTML code that rescales those images such that they all fit inside fixed-size portlet window. This portlet allows the user to select and combine image data generated by other services in order to increase informational density of the dashboard page.

6.5. RRD data analysis and display tool

One of the ideas of the project is to build a solution that offers a formal interface to analyze disjoint data sets. Perhaps the simplest approach to address this issue is to re-use existing tools and their interfaces through a thin layer of templates and known agreement on how infrastructure should instantiate the actual interface invocation. Hence while we do not directly address the data analysis in phase 1 of the project, we nevertheless prototyped an RRD driven service for template based transformation and display of XML time series.

7. Future work

One of the major obstacles in reusing the results of this work is the complexity of technologies based on which proposed solution is built. To overcome this, we now plan to focus on understanding common data integration patterns in order to isolate reusable service components. Ideally the roles of such componentsin a model may be changed by simple rearrangement of the workflow or reconfiguration of the parameters of the service.

8. Conclusions

In this project we have addressed the problem of building monitoring and analysis portals for systems which do not support common semantics for data describing system state. Our solution is based on technologies that allow data integration, transformation, and co-display. In the project we have focused on ease of refactoring and adaptation of data through a specially provisioned integration layer. We have used an XAware engine as the driver for that layer, leveraging its support for rules to implement unification of semantics and specifying the format of the input. For addressing the problem of heuristic correlation of information, we have leveraged a data co-display idea through building experiment dashboards using JBoss-backed portlet technology. In the future we'll focus on allowing more sophistication in how advanced data display tools can be used in conjunction with data generated by user-supplied transformation templates. The purpose of this work is to ensure a better fit of metric data analysis and correlation products to monitoring and troubleshooting requirements of users and facility operators.

References

[MCAS] The Metrics Correlation and Analysis Service project:
http://www.fnal.gov/docs/products/mcas   Accessed on May 14, 2009
[Portlet] The Java Community Process, JSR-000168 Portlet Specification
http://jcp.org/en/jsr/detail?id=168   Accessed on May 14, 2009
[XAware] K Vandersluis 2004 XML-Based Integration with XAware: Unifying Applications and Data in Today's e-Business World *Gulf Breeze, FL: Maximum Press* [JBoss] M Fleury, F Reverbel 2003 The JBoss extensible server *Lecture notes in computer science* Springer
[Pluto] The Apache Reference Implementation of the Java Portlet Specfication
http://portals.apache.org/pluto/index.html   Accessed on May 14, 2009
[eXo] eXo Platform   http://www.exoplatform.com/   Accessed on May 14, 2009
[Mule ESB] Mule Enterprise Service Bus   http://www.mulesource.org/   Accessed on May 14, 2009
[SEDA] Welsh M, Culler D, Brewer E 2001 SEDA: An Architecture for Well-Conditioned, Scalable Internet Services *Proc Eighteeth Symposium on Operating Systems Principles (SOSP18), Chateau Lake Louise, Canada*
[RRD] RRD Tools   http://oss.oetiker.ch/rrdtool/   Accessed May 14, 2009