Collaborative Monitoring and Visualization of HPC Data

Roselyne Tchoua *, Hasan Abbasi *, Scott Klasky *, Qing Liu ", Norbert Podhorszki *,
David Pugmire * , Yuan Tian †, Matthew Wolf ǂ
* *Oak Ridge National Laboratory*, Oak Ridge Tennessee
" *National Institute for Computational Sciences*, Oak Ridge, Tennessee
† *Auburn University*, Auburn, Alabama
ǂ *Georgia Institute of Technology*, Atlanta Georgia

***Abstract*— As simulations begin to scale to extreme processor counts trying to understand the mysteries of the universe, collaboration becomes an essential piece of the scientists' daily life as they work to run, analyze, and process their data from these simulations. Most of the teams that we collaborate with work identically to the way they did in the past, without using effective collaborative tools to share their knowledge with their peers. Collaboration is usually an afterthought, and is often handled in an awkward setting. We believe the best way to introduce collaboration into areas that are resistant by nature is to embed into low level and hidden system components so that scientists collaborate without consciously putting in extra efforts. Based on this hypothesis, this paper presents our work in creating a collaborative system, which allows a diverse set of scientists to work together efficiently. Two of the main aspects of our system are in our use of provenance to associate files and its associated metadata with the information that domain experts are interested in, and an easy-to-use high-performance I/O system which automatically annotates the output file(s) with a unified schema. To accomplish our goals, we leverage an existing I/O framework, ADIOS and an existing web interface, eSiMon, and add new techniques and mechanism to efficiently bring together computation and visualization.**

Keywords-component; collaboration; schema, visualization, data staging, simulation, monitoring, collaboration

I. INTRODUCTION

The workflow for scientists engaged in leadership-scale science has grown quite complicated as the size and complexity of data produced through experiments and simulations has increased many-fold over the last decade. With the prominence of verification and validation at the boundaries of simulation, theoretical and experimental scientists have to work closely together. These researchers are collaborating on a collection of data resources coming from different simulations, experimental apparatuses, and theoretical models. Although computing and collaboration technologies have been incorporated by many of such teams to aid in their workflows, there is still need for the development of additional data management and data collaboration tools. Visualization, analysis, verification, statistical testing, and data format conversion, let alone the "simple" tasks of data transport and storage, are all critical tools in today's leadership scale science efforts. The eSiMon [1] and ADIOS [2] platforms described below represent an approach to flexibly support these sorts of real-life collaborative workflows for scalable science.

As new High End Computing (HEC) platforms continue to push the performance envelope, the complexity of the scientific processes they model and simulate also increases. Effectively utilizing these available computational resources involves more than just acquiring programming skills. Application scientists need to work with experimentalists to validate and verify their data and with computational scientists to use the supercomputers and to visualize the tsunami of data coming out of the simulations.

In this paper we use the example of fusion scientists and describe a new method for better organized collaboration. We aim at facilitating the exchange between different types of scientists and minimizing learning curves as well as time spent outside of their interest. Our goal is to take researchers to the science as fast as possible. We describe our efforts to develop a framework that accommodates this type of collaborative workflows. By adopting a service-oriented approached, the collaboration systems discussed here are structured so that users can dig further into their field of expertise with limited but sufficient knowledge about related and collaborative fields.

II. SYSTEM OVERVIEW AND RELATED WORK

Our experience with application scientists has led us to the conclusion that most simulations do not happen independently. Simulations often require input parameters that represent valid initial physical conditions. To complete a successful run, a scientist needs to generate results comparable or relatable to experimental data. The process of correlating the simulation results to the physical model and experimental results is called Validation and Verification (V&V), which requires a large degree of collaboration between the simulation and the experimental scientists. Visualizing and analyzing the results is yet another subsequent endeavor that required intervention from experts in these fields. While one could envision a situation where these separate groups tackle different components of a problem and periodically converge back to exchange their findings, we believe the solution is an end-to-end service-oriented-approach. The scientific workflow of tasks that lead to knowledge discovery has grown in complexity and size but instead of breaking it into pieces – risking the chance of introducing divergences – we want to keep the underlying structure of the work and develop tools that accommodate it. Therefore, even though we now have teams of scientists solving growing challenges, we want to preserve the method that leads to valid results.

978-1-4673-1382-7/12/$31.00 ©2012 IEEE

There are four main pieces to the architecture as illustrated in Figure 1. The common access point for collaborating scientists is the eSiMon portal which stores both comprehensive and diverse metadata information collected from the simulation and about its users. In other words this information includes details about the data as well as the users' interaction with that data. The sources of the data, either experimental sensors or computational simulations, are two other major parts of the system. In the example of collaborative research in fusion sciences the experiments are tokomak reactors and the simulations (computer programs) run on leadership class supercomputers such as Jaguar at Oak Ridge National Laboratory. The final piece of the system consists of the analysis and visualization workflows that handle the post-processing of the data. eSiMon serves as the glue that connects the disparate data sources as well as the scientific workflows. This significantly reduces the burden on scientists in dealing with the volume of data since eSiMon is designed to both know and understand the multiple pieces and how they interact. In other words the supporting backend to eSiMon is doing the heavy lifting (data lineage, user activity recording and tracking) while the users are presented directly with the scientific variables and phenomena.

Bringing these four components together involves supporting background work that primarily addresses data movement and interpretation. We move the data from the compute nodes to staging nodes [3]; pre-process it to produce preliminary diagnostic images and text tiles to be visualized on eSiMon. eSiMon provides access to experimental data sources that can be used as input files or as measurements for V&V. For data to flow easily from raw data sources to preliminary results we need a common language: a unified data schema. Furthermore, we need continuous and consistent recording of data flow to accurately link different components together. Finally to facilitate use and reuse of the data, we not only record its origin but also users' activity and interactions with the data. If we picture the data essentially moving from raw data to preliminary results, we begin to see eSiMon as a service-oriented one-stop-shop where users can access different stages of the process. Scientists can interact with the data and each other without knowledge of the intricacies encountered and surmounted in the back end.

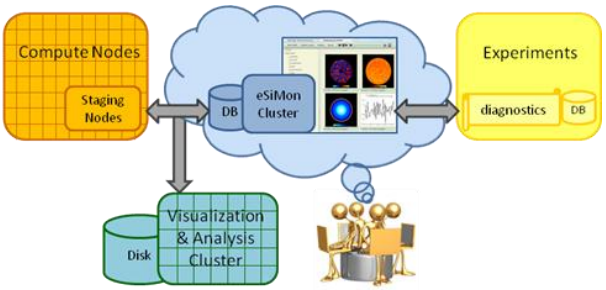

Figure 1. Monitoring HPC data.

There has been extensive research on synchronous and asynchronous collaborative writing and drawing software [4, 5, 6, 7]. The impact of technology on communication between participants has also been explored [8, 9]. Collaboration on large display, touch-screens, white boards and tables [10, 11, 12] has also generated a great deal of interest. Several studies want to improve the way people interact in group meetings where most participants are present. Concurrently, some research in synchronous collaborative visualization [13] that attempts to minimize the effects of time and space between remote collaborators. However, the adoption of such tools can be relatively low. Generally accepted theories suggest the adoption of collaborative tools strongly depends on the extent to which they match the widely varying types of collaboration. With social networking new tools are emerging such as myExperiment [14] that allow scientists to exchange workflow and experiments via the Web. Personalized "dashboards" have also become more popular in an attempt to address the issue of relevance and context for different users [15]. In this paper we identify and address a new and specific type of simulation at the intersection of theoretical and experimental sciences, visualization and analysis.

High performance computing has changed the dynamic of the collaboration in certain scientific teams. There are only a few leadership computing facilities. Access to the resources at these facilities is given according specific metrics and guidelines to local and remote users. Using the supercomputers also requires a level of expertise in computational sciences. This means that even though physical proximity is key to successful collaboration it is very common for teams of simulation scientists to be scattered between experimental laboratories, computing facilities and universities. In addition while scientific visualization of distributed and parallel simulation is often associated with 3D objects, we have found that theoreticians work intensively with 1D and 2D diagnostic graphics when exploring their results and before getting into more in-depth visualization. The work described in this paper does not target "shoulder-to-shoulder" collaboration, nor is it addressing the issues of "live" collaboration between users from different physical locations. Instead, we are interested in the intersections between different fields of sciences and the issues associated with exchanging and handing over work from one team to another. The eSiMon does not require users need to be concurrently be connected. It differs from personalized dashboards since it does not present the same content contextualized for different clusters of users. Rather it emphasizes connections and transitions from one task to another in this end-to-end workflow of scientific tasks. It highlights the gates to different activities as well as the areas where communication is needed before diving further into these activities. How do we embed some of the transitioning knowledge into the tool so that diverse users effortlessly and efficiently find the way into their science and to related fields only when needed? This is where our key contribution lies.

The underlying principle of our work is our SOA approach supported by our metadata-rich I/O, our uniform schema, our provenance tracking system and our staging method. This allows us to embed and record knowledge unbeknown to the users and use it to coordinate some mundane or unfamiliar tasks in the background. eSiMon provides access to different services performed on the simulation or the experimental data. The particular choice, of which tools are used behind the scene, is not necessarily relevant to the user who wants to inspect the

results or perform certain mathematical transformation to the data. Our primary concerns are ease-of-use and performance. More specifically our goal is to seamlessly provide performance.

III. MOTIVATION

Although the techniques, tools and approaches we describe in this paper are broadly applicable to other advanced scientific collaboration environments, we draw particular inspiration from the details of our collaboration with scientists involved in theoretical, experimental, and computational fusion research. We briefly describe here some of the key lesson and exemplars we have learned from those interactions.

A. Experimentation

Extracting valid science from fusion simulations is an involved process starting with the running of the experiments. In our experience with fusion simulation codes, such as XGC [16] and GTC [17], simply preparing the correct inputs for a new simulation run requires considerable effort. The results of a successful run need to be comparable or relatable to experimental data. However, the experimental data obtained from the physical tokamak is dependent on the nature of the diagnostics used. For example, some experiments are designed to measure electron temperature, density pressure profiles and line-integrated density while others measure magnetic fields and fluxes. Due to these distinctions, running a simulation to produce comparable results requires in-depth knowledge of the particular plasma experiment. Selecting the experiment is also a difficult endeavor. As an example, simulation scientists from PPPL agree that the best, and practically the only, way to correctly select an experimental run (typically referred to as shot) from the MDSplus [18] database is by directly speaking with the experimentalist who ran it originally. Since there is very little logging when running the actual experiment, there is no easy way to query the list of experimental shots. For example a user may want to access all shots where the current went over a threshold value. The existing infrastructure does not allow such queries on experimental data. Therefore, the first task at hand for physicists is to speak directly to an experimentalist. This is a significant obstacle to effective collaboration even when the scientists are located at the same institution. With multi-institution, and multi-national, collaborations the complexity of these interaction is ratcheted up to a much greater degree. Moreover because there are no logs on such exchanges, this process is repeated for every single simulation and it is not difficult to imagine how much work is needlessly repeated each time due to the lack of activity tracking and recording.

Once scientists have agreed on an experimental shot, the recorded raw data must be further processed and is not usually ready to use for input and comparison. For example, often this data consists of discrete measurements and does not contain the exact entities that the simulation physicists need. Physicists have to run an additional analysis code to smooth the raw data out and extract the interesting derived variables. It usually takes several such runs to select satisfying initial or equilibrium conditions needed as input to a fusion simulation. Even after the correct experiment is selected and refined, the correct moment in time has to be selected. Since the experiment happens within seconds, pinpointing the best time to serve as input constitutes another level of tuning required in preparation for the simulation. While the scientists keep their own logs of their activities, these logs are often not shared with the broader community and due to the narrow focus of the scientists often result in the occlusion, and even outright loss, of significant insights.

Thus the process of executing a fusion simulation involves detailed discussion between both the simulation scientists and the experimentalists. The actual decisions, as well as the material and knowledge that form the basis of these decisions, are not recorded. Moreover, the benefits are restricted to the scientists directly involved in the simulation and the experiment, despite the interaction being generally valuable for the community. In the end little seems to be done in order to make the preparation for the next run easier or quicker.

B. Simulation

During our experience with the fusion scientists of the CPES [19] project we have developed an I/O framework, a workflow system and a dashboard to help run, monitor and pre-process simulation results for fusion scientists. While physicists can write computer programs that describe their equations, they are usually not computational experts and are not always able to take full advantage of the state of the art computing resources available to them. Therefore they need to work closely with I/O and performance experts in order to produce valid results when using supercomputers. Our data management team has long pursued the goal of assisting HPC users in their quest to get scientific insight from simulations run on the supercomputers. This has been the driving force of our Framework for Integrated End-To-End Technologies for Applications which include three main elements: a fast adaptable I/O (ADIOS), a workflow management system and a collaborative portal (eSiMon). We have helped scientists achieve improved I/O performance. More specifically decisions about where the I/O is done and which resources are used are important to attain satisfactory I/O performance. Furthermore scheduling I/O has considerable impact on these numbers. Physicists are not particularly interested in these issues; they are only concerned with successfully running their simulation. They are aware of known issues: the cost of I/O for HPC application can be substantial and noise from I/O actions can inject undesirable delays into the runtimes of their codes. These inconveniences are tolerable up to point. However as applications scale, they can bring a simulation to a halt and at this point, physicists grow more concerned. Hence, the question that has motivated our work has always been, how do we help application scientists run simulations without having to become I/O experts and how do we facilitate the communication between these two types of experts? The main benefit of ADIOS is that it provides portable, metadata-rich output along with very fast scalable I/O. This allows us to provide sophisticated techniques that take full advantage of the current research and resources without burdening the users. Indeed, with minimum overall knowledge of I/O, physicists describe their variables and

choose an I/O method by simply annotating the ADIOS XML file. Ease-of-use is the key when trying to help pure scientists switch the way they currently work. Offering clear gain in performance is sometimes not enough when it involves too much work from unfamiliar users.

C. *Visualization*

After the experimental data selection and the completion of the simulation run, the next step is to assist our users with exploring the results. Managing and exploring the ever growing amount of data produced on the supercomputers data is a real challenge in itself and once again requires experts in the data management to assist application scientists. The data needs to be moved, reduced, transformed as much as possible before it can be visualized and analyzed. There is currently a gap between visualization experts and physicists. Visualization scientists require certain information and data structures in order to graph data; physicists are not always aware of what is needed for specific software to read and plot their data. The easier or more intuitive way for them to write data out in the code does not necessarily corresponds to the best reading performance for visualization. Also the fact that different scientists use different software further complicates the issue. Once again, without effective communication between these two types of scientists, physicists often end up writing their own personalized routines in whatever visualization tool is available to them for various and arbitrary reasons. This makes the process manual, sometimes cumbersome, not always easy to repeat for the physicists themselves and definitely not easy to share. Some scientists already use self-describing data formats, which helps in ensuring that data written by simulations can be correctly read by visualization tools. Visualization and analysis software tool developers provide "readers" for some of the common data formats, such as HDF5 [20] and NetCDF [21]. Still these file formats do not impose a standard organization for the data of a particular scientific field. Therefore, different scientists may identify the same physical property with a different name, and may utilize a different way of organizing the data. This means that from one scientist to another, one must translate the data organization to interpret it appropriately. Thus the syntactic representation of the name obfuscates the semantic meaning of the data and requires extra steps towards collaboration.

IV. IMPLEMENTATION

Based on our experience with application scientists and their various interactions with other experts we have designed an integrated system.

A. *Lessons Learned*

The HPC community despite its growing size is still a relatively small community which limits our users' studies. On the other hand given a few hundred users on the jaguar computer at ORNL, the fusion scientists we have worked with are a considerable sample of simulation scientists. Based on this experience we gather our lessons learned. A helpful collaborative tool between experimentalists and simulation scientists needs to allow simultaneous and easy-access to supercomputers and experimental data sources. Application scientists and experimentalists should be able to concurrently explore experiments and simulation shots. The system should store and deploy experimentalists' and physicists' custom routines on the data to select the correct input for the simulation. It should allow job submission using previously selected inputs. This should be an environment where simulation results can be easily compared to experimental values for verification and validation. We should integrate with software that promotes and encourages easy note taking and the ability to add comments during the collaboration process. This insures that the work involved is not lost for the future runs. Such information should be exposed to all collaborators and will be especially useful to new users when submitting their own simulations without starting from scratch. Data must be exchangeable and interpretable by different components of the system.

The goal of our collaboration infrastructure is to make this process as transparent and intuitive as possible, regardless of where the experimental data is stored and where the simulation is run. While there are questions about sharing sensitive data, the focus of this paper is on the design and development of the collaboration infrastructure. Development and management of access and sharing policies will be addressed in our future work.

Our approach is geared towards quick-and-easy collaborative online monitoring using data staging services. Key components of our system are: a unified schema, continuous and consistent provenance tracking system, and a data staging method as shown in Figure 2. eSiMon incorporates and gives access to these crucial pieces to a collaborative system.

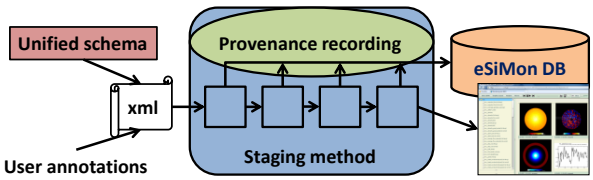

Figure 2. Key System Components.

B. *eSiMon*

eSiMon (electronic Simulation Monitoring) is our one-point-access to every other component of this integrated solution. The key emphasis of this web interface is our persistent emphasis in presenting users with variable names and mathematical function names instead of file and directory names. We use the provenance information and metadata throughout the system to hide IT details from our users. eSiMon provides an "at-a-glance" look into the status and health of the simulation. Multiple collaborators on a team can log on to this web interface and share simulation results. It is our attempt to create a scientific social network for simulation scientists. Team members can log on to eSiMon from any device and browser to view simulation results as they get output from the simulation. Having hooks into different data sources and key components of the framework, eSiMon aims at being an environment where users with general knowledge in the field can access different types of input and output data

regardless of where the data is physically located. Similarly users can execute transformations on simulation data independently of what software and which resources are being used. The type of analysis on eSiMon is light-weight and meant to give initial insight into the wide range of results. It is the first simplified medium between scientists and the simulation. We expect that after using eSiMon users will know exactly where in the simulation they wish to do more-in-depth analysis using their favorite more sophisticated tool.

C. Provenance

Based on the complicated interactions between application scientists and experimentalists, an effective system must have a reliable data lineage component. We have developed a provenance tracking system [22] and have integrated it in our portal eSiMon to record data movement during the simulation. Data provenance tracks the origin and the history of the data, recording along the way its transformation and derivatives. In the past we have specifically focused on this type of information for linking pre-processed images to the original raw data files. This is useful when doing online analysis or downloading sets and subsets of data to a particular machine for further analysis. We now extend this concept to user activities tracking. By establishing a connection between eSiMon and an experimental data source and saving user activities, we record the process one user goes through to select an experimental shot in the Provenance Store. This is not only useful for this user's subsequent runs, but it is especially useful for a collaborator to come along and avoid repeating this work. Moreover, our system that encourages notes and comments so that as the number of run increases in the database we allow queries on past shots. Furthermore by establishing one common access point to experimental data and simulation data, the system provides an environment where users can setup their simulation inputs from experimental data, use them to run their simulations and then compare simulation results with other experimental data. Collaborators using the web portal can see each other's previous runs and simply reuse or tweak the input parameters. Letting the system record all users' activities avoids repeating the exchanges between experimentalists and physicists and skips over the less interesting part of setting up the simulation. This is a great time and effort saver compared to previous methods.

D. Unified Schema and Staging Method

Our solution to the mismatch in semantics between scientists, visualization tools and existing data formats is to develop a standardized schema for HPC data. There has been a lot of effort put into developing a standardized method to exchange data between HPC codes and tools using technologies such XDMF [23]. VizSchema [24] is another similar effort to standardize scientific data by linking common data formats to common software tools. We define our schema as a data model designed for a purpose; in this case mainly visualization and collaboration. Moreover since we embed this schema into state-of-the-art I/O system, we provide performance as an additional benefit. This system insures that data written by the simulation can be visualized independently of user expertise, personal preference or availability of software packages. This decision is more suitable for exchange and collaboration. Physicists do not have to become experts in a specific software package to do their visualization or analysis. As they work in teams, they do not have to impose the use of one single tool for all members since the data and its representation are self-described, and thus portable. This considerably optimizes their communication with I/O and visualization experts.

It is not uncommon to use XML for data models. XML tags and attributes are a convenient way to organize data as long as it is simple enough for human users. Since XML is already an intrinsic part of ADIOS we allow users to annotate their data structure in the ADIOS XML files. These annotations include not only the type and description of data, but also the meshes it is associated with for automatic visualization. The goal here is to, with minimal added information – parts or attributes, allow our staging plotting tool or any other visualization software to interpret the data. Nevertheless, the schema will allow annotating the data in progressively greater detail, by the use of additional attributes. For example more metadata such as units and axis labels can be added to variables subsequently and will be interpreted by the staging method when producing images.

Our standardized schema for the simulation output is comprehensible by popular visualization tools. Furthermore it enables us to create our own reader in ADIOS. ADIOS provides data staging method that utilize supplementary resources to provide additional functionality to the application. Data staging Services use asynchronous methods for data extraction from compute nodes. By moving the data from compute nodes to staging nodes prior to storage, data staging services can considerably reduce the impact of the I/O on the computations. Simultaneously, the available resources in the staging nodes can be used for additional capabilities as part of an I/O pipeline. With this staging method we are able to do pre-processing of the data before it hits storage. Moving the data out of the compute nodes and producing diagnostic files and images on staging nodes during the simulation enables live monitoring of the health of the simulation. This plotting method reads the users' description of meshes and variables using to the new schema and appropriately interpret their data for visualization. Selecting this method like any other ADIOS methods does not require users to change their code. By correctly annotating the XML file, they can choose available I/O methods and take advantage of powerful pre-processing services. When new I/O methods become available, or the visual presentation of a variable is to be changed, ADIOS users do not have to change their code, they simply alter their XML file. While the details of this method seem complex, they key is that users do not have to be I/O experts and understand the details of the method. Below is a sample of an ADIOS XML file describing a variable (pressure) on a 2D mesh (s3d_mesh). The underlined text highlights the user's additions to a typical ADIOS XML file.

```
<var name="xd"  type="integer"/>
<var name="yd"  type="integer"/>
<var name="lx"  type="integer"/>
<var name="ly"  type="integer"/>
<var name="sg1" type="integer"/>
<var name="sg2" type="integer"/>
<global-bounds dimensions="xd,yd"
offsets="sg1,sg2">
     <mesh name="s3d_mesh" type="structured"
ndims="2" nspace="2" />
          <var name="x" type="double"
dimensions="lx,ly "/>
          <var name="y" type="double"
dimensions="lx,ly "/>
     </mesh>
     <var name="pressure" type="double"
dimensions="lx,ly" mesh="s3d_mesh"/>
</global-bounds>
```

Figure 3. Sample of ADIOS XML File schema.

The newly added tags and attributes are processed in ADIOS and the variable pressure automatically gets associated with the s3d_mesh in the binary file. The user no longer has to worry about the best way to write out data or changing their code for visualization tools. Listing the content of the file generated from the XML description above generates the following output.

```
integer /xd              scalar = 512
integer /yd              scalar = 1024
integer /lx              scalar = 24
integer /ly              scalar = 36
integer /sg1             scalar = 3802
integer /sg2             scalar = 3825
double  /pressure {24, 36} = -1/1/0.06/0.73
string  /s3d_mesh/type   attr = "structured"
string  /s3d_mesh/time   attr = "no"
integer /s3d_mesh/ndim   attr = 2
integer /s3d_mesh/dim1   attr = 24
integer /s3d_mesh/dim2   attr = 36
string  /pressure/label  attr = "pressure"
string  /pressure/mesh   attr = "s3d_mesh"
```

Figure 4. Group schema.

A well-defined schema allows reorganizing the data and shuffling variables around without having to write new readers. Automatically generating 1D and 2D diagnostic images allows physicists to directly dive into sections of the output that looks more interesting to them. Currently they manually run series of standard analysis routines, look at those results before focusing on a subset of the data. This is where exciting knowledge discovery really begins. Simulation scientists in our experience have expressed that the work done to dig deeper into a particular result, or to produce a publication-ready figure is something that they do desire to do on their own. This part of the work varies with each simulation. This is where simulation scientists would rather spend their time and effort. The work done at this level will likely not be exactly repeated for subsequent simulations. The scripts used here may not be used again for a long time if ever. In other words the goal is to bring the simulation physicists to this point as efficiently and effortlessly as possible from the preparation of the experimental data, to the automatic generation of preliminary images and then leave them alone to work the way they want but providing linkage and access to all relevant data in one integrated environment.

V. CONCLUSION

The work described in this paper targets the need for effective exchange and collaboration between different types of HPC scientists. To address these needs in the HPC community we designed a uniform standardize schema along with a new staging method for run-time monitoring of the simulation while maintaining persistent records of data lineage. We leverage the existing I/O framework in ADIOS and the existing web interface in eSiMon and add new techniques and mechanisms to efficiently bring together computation and visualization. We treat operations on experimental or simulation data as services offered by eSiMon to teams of scientists. We provide a single access point for all members of diverse sub-areas of a broader field of science; we highlight intersections and areas where exchange and communication are needed; we hide superfluous information and finally we provide performance as a bonus benefit. We envision that in the future, physicists will not only have access to others' data and routines but also be able to make new requests to collaborators. In the future we also plan to extend this abstraction to utilize resource external to the main machine such as smaller scale visualization clusters. The advantage of such an extension is the ability to have longer running computation, custom resources such as GPUs or local storage using fast SSDs. We remain motivated by ease-of-use and user friendliness when assisting users with the visualization of their data to improve and accelerate the process of gaining scientific insight into extreme scale data sets.

ACKNOWLEDGMENT

This research and development has been supported in part by the National Center for Computational Sciences at the Oak Ridge National Laboratory, by DOE SCIDAC-2: Scientific Data Management Center for Enabling Technologies grants DE-FC02-01ER25486 and DE-FC02-ER25809, and by SciDAC Fusion Simulation Prototype Center for Plasma Edge Simulation.

REFERENCES

[1] Electronic Simulation Monitoring (eSiMon): http://www.olcf.ornl.gov/center-projects/esimmon/.

[2] J. Lofstead, K. Schwan, S. Klasky, N. Podhorszki, and C. Jin, "Flexible io and integration for scientific codes through the adaptable io system (adios)," In Challenges of Large Applications in Distributed Environments (CLADE), 2008.

[3] Hasan Abbasi, Matthew Wolf, Greg Eisenhauer, Scott Klasky, Karsten Schwan, Fang Zheng, "DataStager: Scalable Data Staging Services for Petascale Applications," Proceedings of HPDC 2009.

[4] R. M. Baecker, D. Natos,I. R. Posner and K. L. Mawby, "The User-centred Iterative Design of Collaborative Writing Software," In Proceedings, InterCHI'93, pp. 399-405, 1993.

[5] A. Karsenty, C. Tronche, and M. Beaudouin-Lafon. GroupDesign: shared editing in a heterogenous environment. USENIX Journal of Computing Systems, 6(2) pp. 167-195, 1993.

[6] R. Fish, R. Kraut and M. Leland, "Quilt: a collaborative tool for cooperative writing," In Proceedinsg of Conference on Office Information Systems, 1988.

[7] B. Lewis, J. Hodges, “Shared Books: Collaborative Publication Management for an Office Information System,” in Proceedings of the ACM Conference on Office Automation Systems, pp. 197-204, 1988.

[8] J. Galegher, and R. Kraut. Technology for intellectual teamwork: Perspectives on research and design. In J. Galegher, R. Kraut, & C. Egido (Eds.) Intellectual Teamwork (pp. 1-20). Hillsdale, NJ: Lawrence Erlbaum Associates. 1990.

[9] J. S. Olson, M. Storrosten and M. Carter, “How a Group-Editor Changes the Character of a Design Meeting as well as its Outcome,” in Proceedings of CSCW’90, pp. 91-98. 1990.

[10] S. Erold, R. Bruce, D. Goldberg, F. Halasz, W. Janssen, D. Lee, K. McCall, E. Pedersen, K. Pier, J. Tang, B. Welch. Liveboard: “A Large Interactive Display Supporting Group Meetings, Presentations and Remote Collaboration,” in Proceedings of the ACM CHI’92 Conference on Human Factors in Computing Systems, pp. 599-607. 1992.

[11] I. Tollinger, M. McCurdy, A. H. Vera, P. Tollinger, “Collaborative Knowledge Management Supporting Mars Mission Scientists,” in Proceedings of the 2004 ACM conference on Computer supported cooperative work CSCW 04 (2004)

[12] T. Prante, N. Streitz, P. Tangler. Roomware, “Computers disappear and Interaction Evoles,” in Computer 37, 12, pp. 47-54. 2004.

[13] S. Cutchin, C. Bajaj, "Web based collaborative visualization of distributed and parallel simulation," pvg, pp.47-54, in Proceedings of the 1999 IEEE symposium on Parallel visualization and graphics (PVGS '99), 1999.

[14] D. De Roure, C. Goble, R. Stevens. Designing the myExperiment Virtual Research Environment for the Social Sharing of Workflows. Proceedings of the Third IEEE International Conference on eScience and Grid Computing, pp. 603-610. eScience 2007.

[15] T. Palpanas,P. Chowdhary, G. A. Mihaila, F. Pinel, "Integrated modeldriven dashboard development." Journal of Information Systems Frontiers, 9(2-3). Jul 2007.

[16] S. Ku, C. Chang, M. Adams, J. Cummings, F. Hinton, D. Keyes, S. Klasky, W. Lee, Z. Lin, S. Parker and the CPES team, “Gyrokinetic particle simulation of neoclassical transport in the pedestal/scrape-off region of a tokamak plasma,” Institute of physics Publishing Journal of Physics: Conference Series, Vol. 46, pp. 87–91, 2006.

[17] L. Oliker, J. Carter, michael Wehner, A. Canning, S. Ethier, A. Mirin, G. Bala, D. Parks, Patrick Worley Shigemune Kitawaki, and Y. Tsuda. Leading computational methods on scalar and vector hec platforms. In Proceedings of SuperComputing 2005, 2005.

[18] MDSplus: http://www.mdsplus.org/index.php/Introduction.

[19] Center for Plasma Edge Simulation (CPES). http://www.scidac.gov/FES/FES_CPES.html.

[20] Hierarchical Data Format (HDF5) Group: http://www.hdfgroup.org/HDF5/.

[21] Network Common Data Form (NetCDF): http://www.unidata.ucar.edu/software/netcdf/.

[22] P. Mouallem, R. Barreto, S. Klasky, N. Podhorszki, M. Vouk, “Tracking Files using the Kepler Provenance Framework,” 21st SSDBM conference. (New Orleans, LA, June 2009).

[23] eXtensible Data Model and Format (XDMF): http://www.xdmf.org/index.php/Main_Page.

[24] S. Shasharina, J. R. Cary, S. Veitzer, P. Hamill, S. Kruger, M. Durant, and D. A. Alexander, VizSchema - Visualization Interface for Scientific Data, IADIS International Conference, Computer Graphics, Visualization, Computer Vision and Image Processing, 2009, p. 49.