

Lawrence Berkeley National Laboratory

Scientific Data

Title

MIQS

Permalink

<https://escholarship.org/uc/item/3zj0q27z>

ISBN

9781450362290

Authors

Zhang, Wei
Byna, Suren
Tang, Houjun
et al.

Publication Date

2019-11-17

DOI

10.1145/3295500.3356146

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

Wei Zhang
Texas Tech University
Lubbock, Texas
X-Spirit.zhang@ttu.edu

Suren Byna
Lawrence Berkeley National
Laboratory
Berkeley, California
sbyna@lbl.gov

Houjun Tang
Lawrence Berkeley National
Laboratory
Berkeley, California
htang4@lbl.gov

Brody Williams
Texas Tech University
Lubbock, Texas
Brody.Williams@ttu.edu

Yong Chen
Texas Tech University
Lubbock, Texas
yong.chen@ttu.edu

ABSTRACT

Scientific applications often store datasets in self-describing data file formats, such as HDF5 and netCDF. Regrettably, to efficiently search the metadata within these files remains challenging due to the sheer size of the datasets. Existing solutions extract the metadata and store it in external database management systems (DBMS) to locate desired data. However, this practice introduces significant overhead and complexity in extraction and querying. In this research, we propose a novel Metadata Indexing and Querying Service (MIQS), which removes the external DBMS and utilizes in-memory index to achieve efficient metadata searching. MIQS follows the self-contained data management paradigm and provides portable and schema-free metadata indexing and querying functionalities for self-describing file formats. We have evaluated MIQS with the state-of-the-art MongoDB-based metadata indexing solution. MIQS achieved up to 99% time reduction in index construction and up to 172k \times search performance improvement with up to 75% reduction in memory footprint.

CCS CONCEPTS

• **Computing methodologies** \rightarrow **Parallel computing methodologies**.

KEYWORDS

Metadata Search, HDF5 Metadata Management

ACM Reference Format:

Wei Zhang, Suren Byna, Houjun Tang, Brody Williams, and Yong Chen. 2019. MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats. In *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '19)*, November 17–22, 2019, Denver, CO, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3295500.3356146>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SC '19, November 17–22, 2019, Denver, CO, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6229-0/19/11...\$15.00
<https://doi.org/10.1145/3295500.3356146>

1 INTRODUCTION

Large-scale scientific applications, including experiments, observations, and simulations, generate enormous amounts of data. [3, 4, 8, 11, 24, 25, 27, 31]. Due to upcoming construction of even larger scientific infrastructures and the increasing demand on the fine nature of the data, the volume of scientific data is expected to grow rapidly. Such growth imposes a substantial challenge to scientists - to survive in the ocean of data, and particularly, to efficiently find the data they require [38, 40, 43, 46, 48].

Self-describing data and file formats, such as HDF5 [13], netCDF [33], ADIOS-BP [26], and ASDF [14], are extensively used in scientific applications. In these file formats, the metadata of each data object is stored alongside the data object. This practice makes these formats both self-describing and self-contained. For self-describing and self-contained data management paradigm, it is ideal that the metadata search functionality is included as part of the self-describing and self-contained data management solution. However, most self-describing and self-contained data management solutions do not provide such internal metadata search functionality. Instead, finding interesting data in multiple files is achieved by iterating through all the files and performing pattern matching on the metadata one by one. Such a procedure is easy to implement and works well when there are a limited number of files and the amount of metadata is small. However, when the number of data objects in a file is very large, sifting through the metadata objects within one file becomes time-consuming. Many applications further compound this issue by generating a large number of data files. For instance, cosmology observations store the images of the sky either hourly or daily and stores all the datasets into multiple self-describing data files [24].

While metadata management and indexing have been well-studied in the context of file systems [23, 28, 38, 41, 47, 48], metadata indexing for self-describing data formats remains underdeveloped. Due to the lack of effective metadata search capabilities provided by the runtime libraries and tool-chains packaged with these self-describing data formats, scientists often use a separate database management system (DBMS) to facilitate the process of finding the required data. Typical examples can be seen from BIMM [20], EMPRESS [21], the SPOT Suite [9], and JAMO [18], where relational databases (e.g. SQLite [39], PostgreSQL [32]) and NoSQL databases

(e.g. MongoDB [29]) are used for maintaining the metadata and providing the metadata search functionality needed for locating the required data from the self-describing data files.

However, a database solution disjoint from the self-describing data files is incompatible with the self-describing and self-contained data management paradigm. This approach requires extra overhead in deployment as well as constant maintenance in order to keep the database updated. In addition, initial deployment of the database system for metadata search necessitates reading the metadata from the self-describing files and then loading it into the database. This duplicates the metadata at two places and hence leads to storage redundancy. Moreover, the database has to be either made available or migrated when the self-describing files are transferred to a different site, which is a complicated process, not to mention that the entire database has to be updated with the new location of file paths.

Toward the goal of providing a self-contained metadata search service for self-describing file formats, we propose *MIQS*, a novel Metadata Indexing and Querying Service. MIQS is designed to be a library that can be integrated into the existing self-describing and self-contained data management solutions. It provides metadata indexing and querying functionality for self-describing data formats in a way that complies with the self-describing and self-contained data management paradigm. By introducing an in-memory index structure, along with an on-disk persistence mechanism of the index, MIQS provides the capability to build portable indexes and an efficient metadata query service. Applications that utilize MIQS can perform metadata search without consulting an external DBMS. This philosophy complies with the self-describing and self-contained data management paradigm.

We have developed a prototype implementation of MIQS to support the HDF5 library and to study its effectiveness. We have conducted evaluation tests on the Edison supercomputer hosted at the National Energy Research Scientific Computing Center (NERSC). Our experiments against over 144 million attributes from 1.5 million objects in 100 astrophysics data files [24] show that, compared to a state-of-the-art MongoDB powered indexing solution, MIQS achieved up to 99% time reduction in index construction and searched metadata about 172k× faster. Moreover, MIQS used only up to 25% of the overall storage footprint that the MongoDB-based solution used for these datasets.

The contributions of this research are summarized as follows:

- We identify the drawbacks of existing DBMS-powered metadata indexing and querying solutions for self-describing file formats.
- Following the principle of self-describing and self-contained data management paradigm, we introduce an integrated, lightweight metadata indexing and querying service for self-describing data files. The in-memory indexing data structure achieves significantly better query performance without transforming or duplicating metadata, as compared to existing database-powered solutions. The index persistence mechanism also makes the index reconstruction easier and ensures the portability of the metadata search service.
- We develop a prototype implementation of MIQS to support the HDF5 library and conduct extensive evaluations to

validate the design. The evaluations confirm that MIQS is efficient in building metadata indexes and offers better query performance than database-powered solutions, with less storage consumption. Additionally, MIQS is both portable and transparent in nature, which promotes the design philosophy of the self-describing and self-contained data management solution.

The rest of this paper is organized as follows. In Section 2, we review self-describing data formats such as HDF5 and existing solutions for metadata indexing and searching. We introduce the design of MIQS in detail in Section 3. In Section 4, we present the experimental evaluation results of MIQS. We conclude this research and discuss future work in Section 5.

2 BACKGROUND AND RELATED WORK

In this section, we briefly review self-describing data formats, using HDF5 as an example. We also discuss existing research and solutions that provide metadata search functionality for self-describing data files.

2.1 Self-Describing Data Formats and HDF5

Self-describing data formats, such as HDF5, netCDF, ADIOS-BP and ASDF, are designed to provide one-stop data management solution with no dependency on other data management solutions. In these data formats, the metadata is stored alongside the data itself, providing description, interpretation, and even definition of the data objects. This fusion between data and its metadata allows the self-describing formats to provide users with programming interfaces and tools that facilitate portable one-stop data management solutions. Applications that adopt these APIs and toolchains can be expected to behave in the same way regardless of where the applications are deployed. Data analysis performed on these self-describing files need not consult other data management solutions as the descriptive information is already contained in the data formats themselves.

HDF5 (Hierarchical Data Format version 5), a typical example of the self-describing data format, is one of the most frequently used data formats in scientific fields, which a large number of users rely on for scientific data management and exchanges [16]. When considering the case where we have a set of HDF5 files managed by the file system, there is a deep hierarchy across both the file system and also the HDF5 data format, as shown in Figure 1.

In a file system (e.g. GPFS [34] or Lustre [35]), each directory may contain sub-directories and also the HDF5 files. Within each HDF5 file, there are two types of objects - groups and datasets. The groups function as internal nodes within the hierarchy while the datasets play the role of leaf nodes. Attached alongside each of these data objects (either a group or dataset) are metadata attributes. Different objects may share the same set of attributes, but the value of these shared attributes can be different from one another. The hierarchical organization of these data objects introduces attribute inheritance between parental data objects and child data objects. In other words, the metadata attribute values of a group also apply to all child groups and datasets underneath.

Not all self-describing data formats follow the hierarchical organization. For instance, netCDF has a flat organization of all variables.

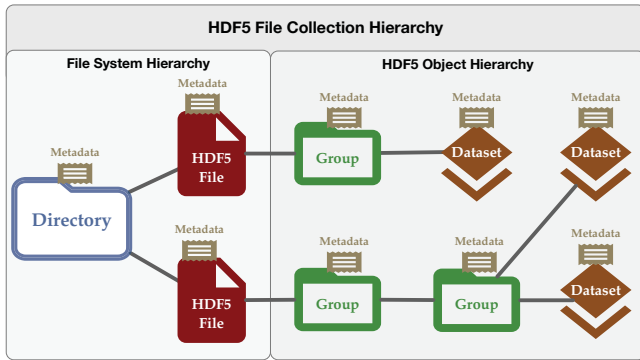


Figure 1: HDF5 file collection hierarchy

However, flat organization is a special case of hierarchical organization, where there is only one level in the hierarchy. As such, we consider metadata indexing and querying using a hierarchical organization to be generalizable to all self-describing data formats. In the rest of this paper, our discussion of metadata indexing and querying will focus on hierarchical organization formats such as HDF5.

2.2 Metadata Indexing and Querying over Self-Describing Files

The metadata in self-describing data formats provides users with descriptive information about the underlying dataset and therefore is prevalently used for finding required data. In self-describing data files, the metadata can be seen as a collection of attributes. Each attribute can be represented as a tuple $\langle k, v \rangle$, where k represents the attribute name and v represents the attribute value. Finding required data can then be accomplished by issuing metadata queries that utilize these key-value pairs. These queries may search for the identifiers of data files or data objects that match given query. Each query contains query target \mathcal{T} and query condition $Q = \langle q_k, q_v \rangle$, where q_k and q_v are the attribute name and attribute value, respectively, in the query condition.

Table 1: Structure of metadata queries with a query target and a query condition

| \mathcal{T} | $Q = \langle q_k, q_v \rangle$ |
|---------------|--|
| Data files | $\langle \text{BESTEXP}, 113919 \rangle$ |
| Data objects | $\langle \text{AUTHOR}, \text{John} \rangle$ |

According to the Digital Curation Conference (DCC) [10], the major data types of metadata attributes are strings and numbers. In particular, the attribute names are of strings and attribute values are either strings or numbers. Therefore, in this study, we focus on metadata queries with q_k to be a string and q_v to be a string or a number. For example, as shown in Table 1, one metadata query may ask for the identifiers of data files that contain attribute key-value pair $\langle \text{BESTEXP}, 113919 \rangle$, where attribute name “BESTEXP” is a string and its value “113919” is a number. Another metadata query may look up the identifiers of the data objects with attribute

key-value pair $\langle \text{AUTHOR}, \text{John} \rangle$, where “AUTHOR” is a string and “John” is also a string.

Most self-describing data formats come with libraries and tools that enable access to the metadata of each data object in a single file. However, many contemporary applications tend to store data into multiple self-describing files instead of a single file [24] in order to avoid performance issues and to be able to better cope with different file systems. In such scenarios, even with the libraries and tools provided by the self-describing data formats, finding the required data over a large number of self-describing files remains a challenging task. Due to the lack of libraries to provide metadata search capabilities over a collection of multiple self-describing files, scientists have to build external metadata search system on their own with the help of a database management system (DBMS).

Examples of such DBMS-powered approaches are BIMM [20], EMPRESS 2.0 [21], JAMO [18], and SPOT Suite [9]. In BIMM [20] and EMPRESS 2.0 [21], the metadata is stored in relational database management systems (RDBMS), such as PostgreSQL [32] and SQLite [39]. It is difficult, however, to accurately reflect the hierarchical organization of metadata for self-describing files using the tabular model of RDBMS. Some other solutions, such as JAMO [18] and SPOT Suite [9], use document database (MongoDB [29] in particular) to avoid the complexity and effort of transforming metadata into tabular format. In these MongoDB-powered solutions, metadata is transformed into BSON format [7] (the binary format of JSON [19]) first, and the BSON representation of the metadata is then stored in MongoDB for efficient metadata search. From the BSON documents in MongoDB, MongoDB indexes have to be created on specified attributes to accelerate queries against these attributes.

However, all database-powered solutions share common drawbacks. In order to explain these drawbacks, Figure 2 provides an abstract view of these existing DBMS-powered solutions. As shown in this figure, the self-describing data files are stored on a shared file system, and applications running on the HPC system can access the data and metadata in these files concurrently. Apart from the HPC system, an external, dedicated database system is built to facilitate metadata search. The metadata of the self-describing data files is first transformed into another format that complies with the data model of the database system, and is then imported into the database system. This operation duplicates the metadata of the self-describing data files and stores them into the external DBMS. To search metadata, applications on the HPC system have to issue queries through the client-side APIs of the DBMS. In this case, metadata queries and the querying results are transmitted over the network between the HPC system and the database. Such a solution disregards the self-containing property of self-describing data formats by introducing external data source for metadata management and hence suffers from the following drawbacks:

- **Performance issue:** Since applications need to communicate with the external database over the network, the communication overhead can be significant. Particularly, when applications have a high degree of concurrency, metadata search queries can overwhelm or throttle the network.
- **Storage redundancy:** Since the self-describing data files already contain the metadata, duplicating the metadata in

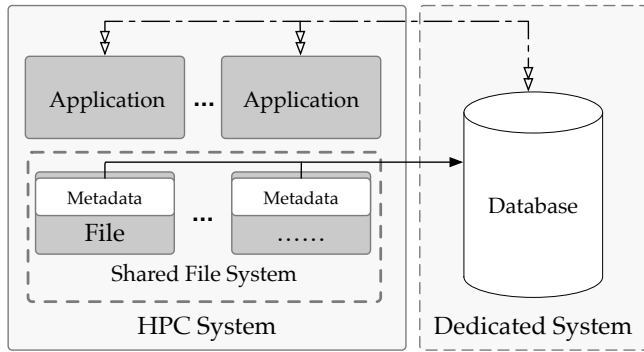


Figure 2: A high-level view of database-powered metadata search solution

any form in DBMS can introduce storage redundancy for storing an extra copy of the metadata.

- Data model adaption:** The data model in a database is different from that of the metadata in self-describing files. In the aforementioned DBMS-powered solutions, the metadata in self-describing files has to be reformatted into either tables in RDBMS or BSON documents in MongoDB, which requires extra efforts and time in transforming data formats.
- Schema constraints:** The data model in RDBMS or document-based databases require the user to know about the data schema used in these systems. For example, in a RDBMS, users have to know on which table an index should be built and from which table the required data records can be retrieved from. In MongoDB, the metadata is represented as BSON document where BSON objects follow the same hierarchical organization found in the original metadata. Therefore, when building index for the attributes, the users must have knowledge on the data schema of the BSON document and they have to know the paths to the data objects where the attributes belong to. As such, when issuing a metadata query in MongoDB, it is necessary for the users to specify the path to the attribute in the query condition. For example, if attribute “BESTEXP” belongs to both object “/3540/55127” and object “/3540/55127/100/coadd”, users have to explicitly indicate two BSON paths for building the index on “BESTEXP”: “3540.55127.BESTEXP” and “3540.55127.100.coadd.BES-TEXP”. Accordingly, when performing a query for “BESTEXP”, users still have to specify the path of the object containing “BESTEXP” [30]. However, when building the metadata index and issuing metadata queries, users are typically aware of only the attribute name (e.g. “BESTEXP”) and attribute value (e.g. 113919). The hierarchical schema of the metadata is hidden from the users and the only way to know the path to the owner object of the attribute is to scan the entire dataset.
- Maintenance demand:** The database system runs on top of the operating system. Consequently, any update to the database system or operating system can cause compatibility issues, which require non-trivial efforts in rectifying them.

- Portability and mobility:** The external database and the metadata search solution are not seamlessly integrated into the self-describing data formats, which introduces complexity in deployment, and significantly affects the mobility of the data files and the portability of the metadata search solution.

3 METHODOLOGY

In this section, we introduce MIQS - a novel metadata indexing and querying service for self-describing data formats. We first present an overview of the MIQS design, and then introduce the index construction procedures of MIQS. Afterwards, we introduce how MIQS serves metadata queries.

3.1 Overview of MIQS Design

MIQS is designed to be a software library providing metadata indexing and querying service that complies with the self-describing and self-contained data management paradigm and can be integrated into the querying APIs or tools of existing self-describing data formats such as HDF5 querying module [15]. Without duplicating the metadata that is already included in the self-describing data files, it only maintains the index of the metadata to facilitate efficient metadata query. As shown in Figure 3, given a collection of self-describing files, MIQS allows each process to create and to maintain a copy of the in-memory index on specified attributes after scanning the data files. Since the index will maintain the relationship between attribute key-value pairs and file/object paths, even a full copy of the index will not consume a significant amount of memory. Once the in-memory index is built, processes can access in-memory index via the metadata search service provided by MIQS. In other words, each process is independent and self-contained. Different from some existing distributed metadata management solutions [17, 40, 42, 45, 47], MIQS follows a shared-nothing parallel architecture which ensures efficient and independent direct memory access for metadata search without network operations.

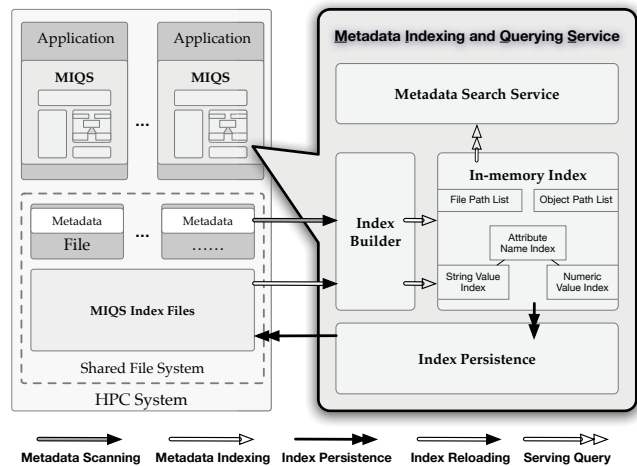


Figure 3: Overview of the architecture of MIQS-based solution

Processes with MIQS can also store the in-memory index to on-disk index files via the index persistence mechanism of MIQS. The MIQS on-disk file can be included as part of the self-describing file collection and can be transferred to any place alongside the self-describing data files. In addition, when transferring a self-describing dataset to a different HPC site, MIQS can quickly reload its on-disk index files into memory to reconstruct in-memory index.

This design eliminates the need for a dedicated external database system. Instead, metadata index files are maintained alongside the datasets while the in-memory index is integrated into existing solutions, which are completely self-describing and self-contained. As a consequence, improved portability and mobility are achieved. Also, in this design, there is no need to maintain databases or to know about data schema in order to use the metadata search service. This design also avoids the effort in adapting data models. Moreover, it does not duplicate the metadata in self-describing data files into other data sources like DBMS, which avoids storage redundancy. The search performance can also be expected to improve since there is no network communication between applications and an external database system.

As the datasets in most scientific applications are generated once and will be read many times for analysis without change [6, 12, 37], our current design does not yet support index updates if a self-describing dataset is modified while being searched. In such a scenario, MIQS needs to reconstruct the indexes to continue performing metadata searches against a modified dataset.

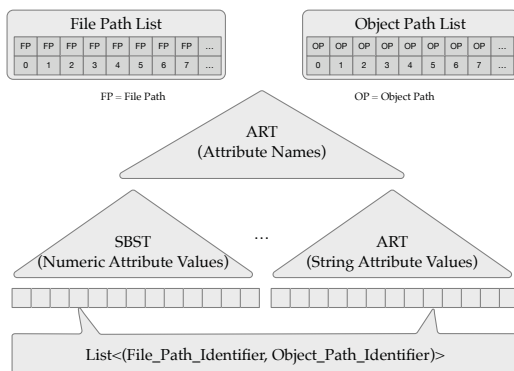


Figure 4: MIQS index data structures

3.1.1 *In-memory Index.* As shown in Figure 4, in MIQS, the in-memory index contains:

- (1) Two path lists serving as mapping tables between path strings and integer identifiers. We use this data structure to store both global file path and global object path mapping tables. In this study, we use the tagged linked list from libhl [2] to implement this functionality.
- (2) A root-level adaptive radix tree (ART) [22] for the attribute keys (name strings).
- (3) A secondary tree-based index at each leaf node of the root-level ART. The secondary indexes are designed for indexing attribute values using a self-balancing search tree (SBST) [36] for numeric values and an ART for string values. (In MIQS,

we use the red-black tree in libhl [2] for the implementation of SBST.)

- (4) A list containing multiple tuples at each leaf node of each secondary index where each tuple contains the file path identifier and the object path identifier. These identifiers correspond to those found in the global file path lists and global object path list which helps avoid storing space-consuming strings, i.e. the repetitive file and object paths.

This combination of ARTs and SBSTs is able to support exact queries on attribute names, string attribute values and numerical attribute values. Also, it opens the possibility for prefix/suffix queries on attribute names and string attribute values as well as range query on numbers. In this study, our major contribution is to eliminate 3rd-party data sources and maintain the property of the self-contained data management paradigm. Therefore, we only discuss exact queries on strings and numbers in order to showcase the benefit of the self-contained and portable design of MIQS. Furthermore, as a result of the direct access to the local in-memory index, both index construction and metadata search are expected to be much more efficient than database-powered solutions.

3.1.2 *Index Persistence Mechanism.* The MIQS index files are designed to permanently store indexes, alongside the corresponding self-describing datasets. An index file can then be loaded into memory quickly for metadata searches when an application uses the dataset. It can also be used to help with the recovery of the metadata querying service if the application process is terminated. These index files can be transferred along with the self-describing data files, so that applications integrated with MIQS can perform efficient metadata queries with minimal effort devoted to setting up the metadata search service.

3.1.3 *Schema-Free Metadata Index Model.* As mentioned in Section 2.2, in the DBMS-powered solutions, the indexing and querying procedures always require information about the data schema. In MIQS, when building the index, users only have to specify the name of the attributes without knowing the host of the attributes. MIQS will match every attribute with the specified name, regardless of the hierarchical schema of the metadata. Also, for metadata search, users only need to be aware of the type of query target, the query attribute name q_k and the query attribute value q_v in the query condition. Users do not need to know the specific structures of the metadata. There is also no need to support a full-fledged SQL-like query language.

3.1.4 *Portability and Mobility.* As discussed in 3.1.2, MIQS stores the index files alongside the self-describing files. Consequently, whenever the combined self-describing files and index files arrive at a scientific computing facility, and the metadata indexing and querying service can be immediately put to use once the dataset is prepared and the application is deployed. This practice conforms to the self-contained data management paradigm. Therefore, we consider the portability and mobility of the MIQS solution for self-describing files to be favorable when compared to DBMS-powered metadata search solutions.

3.2 Index Construction

In MIQS, there are two types of index construction procedures. The first is to build the index from scratch, which involves metadata scanning (optional), metadata indexing and index persistence. We call this procedure “initial indexing”. The other is called “index reloading”. As shown in Figure 3, for initial indexing, the MIQS index builder has to retrieve metadata attributes from the self-describing data files. This is done by scanning the metadata in an existing set of files or by receiving the attributes in an in-situ fashion when the data files are generated. When an attribute is encountered, the index builder retrieves the attribute name and value, along with the corresponding object and file path and creates an index in the memory. After the metadata scanning and indexing are finished, MIQS stores the in-memory index into the MIQS index files using its index persistence mechanism. For index reloading, the MIQS index builder can load all the MIQS index files to rebuild the in-memory index. Our compact index file format is designed to support rapid index recovery. Next, we will describe the initial indexing procedure in detail.

3.3 Initial Indexing

When MIQS index files do not exist yet, we perform initial indexing procedure. The initial indexing procedure includes three sub-procedures, which are metadata scanning (optional), metadata indexing, and index persistence.

3.3.1 Metadata Scanning. When building metadata index on an existing set of data files for the first time, MIQS will perform metadata scanning procedure. Since MIQS is designed for applications running in a parallel environment on HPC systems, MIQS leverages parallelism starting from the index construction phase. Consider a parallel application with n processes where process rank r ranges from 0 to $n-1$. In this scenario, each process maintains a file counter, given as *file_counter*, and that is incremented each time a file is encountered during metadata scanning. The following equation can be used to determine whether or not a given file should be scanned:

$$p = \text{file_counter} \% n \quad (1)$$

When $p = r$, process r will perform metadata scanning on the encountered file; otherwise, it will skip this file and continue. By utilizing this mechanism, MIQS is able to leverage data parallelism across different data files and to avoid the I/O contention occurred when multiple processes are accessing the same data file concurrently.

During the metadata scanning procedure, the MIQS index builder scans four elements from a self-describing file collection (as shown in Figure 5). These elements include :

- The directory that contains the self-describing files. We call this directory the root directory of the data file collection (we use root directory for this meaning in the rest of this paper).
- The self-describing data files.
- The objects in the self-describing files, including groups and datasets.
- The metadata attributes attached to each data object in the self-describing data file.

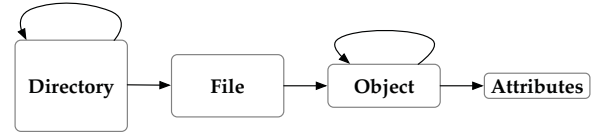


Figure 5: Four elements in MIQS metadata scanning. The arrow shows inclusive relationship.

For scanning metadata, MIQS defines a series of operations on the aforementioned four elements. All these procedures take a global memory pointer called *index_anchor* as a parameter. This memory pointer references the address of the MIQS index root. As part of each procedure, MIQS recursively scans the files of each sub-directory within the root directory.

For each file scanned, MIQS invokes the *ONFILE* procedure. For every object within a file, the *ONOBJECT* is called with the file path and the object path as parameters.

In the *ONOBJECT* procedure, MIQS extracts the metadata attributes for the encountered data object and prepares the index record for each attribute. Since multiple attributes may share the same set of file and object paths, if MIQS stores the path strings alongside each attribute key-value pair, these path strings will be repeated among index records of different attributes for many times. The repetition of path strings will lead to significant memory overhead. Thus, for each attribute to be indexed, MIQS takes the related file and object paths, puts them into the file path list and object path list, respectively, and records their integer identifiers. Afterwards, both identifiers are combined into a tuple $t = \{\text{file_path_identifier}, \text{object_path_identifier}\}$. This design replaces space-consuming path strings that are of variable lengths with fixed-size and space-saving integers. Based on the key and value of the attributes, MIQS is then able to build an index for the attribute with tuple t by calling the *CREATEINDEX* procedure. When finished, MIQS recursively scans and invokes the *ONOBJECT* procedure on any existing child objects.

3.3.2 Metadata Indexing. Given the metadata attribute name *key* and attribute value *value*, along with the tuple t storing the path identifiers of the file and the data object which the metadata attribute comes from, MIQS performs the metadata indexing procedure. As shown in Algorithm 1, in the *CREATEINDEX* procedure, MIQS first extracts the root ART from the index anchor (which is a global pointer to the in-memory index) and then inserts the attribute name into it. By doing so, MIQS is able to locate the corresponding leaf node of the attribute name at the root-level adaptive radix tree (a.k.a ART) where a self-balancing search tree (a.k.a SBST) or an ART places its root. If it is a numeric value, MIQS inserts the value into the SBST and inserts the tuple t into the list which is linked to the corresponding leaf node. Likewise, if it is a string, MIQS inserts it into the ART and links the tuple with the corresponding leaf node of the ART.

After the metadata indexing procedure, each process of the parallel application maintains partial indexes for approximately $1/n$ of the metadata in the entire self-describing file collection. The index persistence mechanism ensure that each process gets a full copy of the index for the entire file collection.

Algorithm 1 Create Index

```

1: procedure CREATEINDEX(key, value, t, index_anchor)
2:    $ART_{root} \leftarrow index\_anchor$ 
3:    $leaf_k \leftarrow insert\_into\_ART(key, ART_{root})$ 
4:   if value is a string then
5:      $ART_{value} \leftarrow get\_ART(leaf_k)$ .
6:      $leaf_v \leftarrow insert\_into\_ART(value, ART_{value})$ .
7:      $link(leaf_v, t)$ .
8:   else if value is a number then
9:      $SBST_{value} \leftarrow get\_SBST(leaf_k)$ .
10:     $leaf_v \leftarrow insert\_into\_SBST(value, SBST_{value})$ 
11:     $link(leaf_v, t)$ .
12:   end if
13: end procedure

```

3.3.3 *Index Persistence.* Once the MIQS index builder finishes metadata indexing for the metadata in all data files, the MIQS index persistence mechanism is invoked and the in-memory metadata index is stored onto disk.

Figure 6 shows the layout of MIQS compact index file format and how the in-memory index is stored in such a format. When

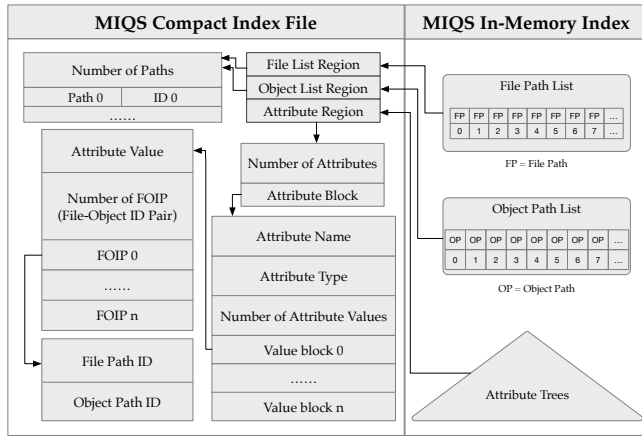


Figure 6: MIQS compact index file format

persisting the in-memory index, MIQS first writes the file path list and object path list in succession. Each path list is persisted as a frame that starts with a frame header. The frame header contains an integer number that represents the number of items in the path list. After the frame header, the items in the path list are stored as a series of item blocks. Each block contains the path string and also the numeric identifier.

After file list frame and object list frame, there is the attribute frame. The attribute frame also starts with a frame header containing the total number of attributes that are maintained in the in-memory index. After the frame header, there are a series of attribute blocks.

The attribute block starts with the attribute header containing the attribute name, the data type of the attribute and the number of unique attribute values. We differentiate three data types supported in MIQS: an integer value type (including all sizes of integers), a

floating-point value type, and a string type. Since each attribute may have different attribute values of the same data type, MIQS records the number of unique attribute values at the end of the attribute header and then write a series of attribute value blocks (or value blocks for short) after it.

Each value block starts with the actual attribute value, the number of matched File-Object ID pairs (FOIPs) - the tuples created in ONOBJECT procedure, and then the actual FOIP containing the file path identifier and the object path identifier.

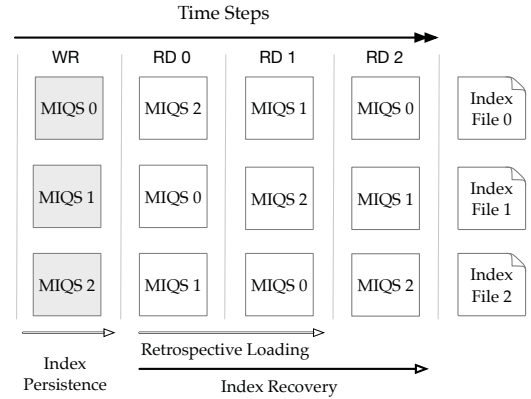


Figure 7: MIQS parallel index file access. “WR” stands for write operation, “RD” stands for read operation.

Following the above procedure, each process will write its own in-memory index into a separate index file with the process rank in the index file name (shown as “WR” step in Figure 7). Thus far, the index persistence procedure is finished.

3.4 Index Loading

3.4.1 *Retrospective Loading.* After the index persistence procedure, each process only maintains the index of the files it has scanned. In order to have a full copy of the index of the entire file collection, MIQS performs “retrospective loading” procedure. For an application with n processes, after these processes finish their own index persistence procedure, each process $r(0 \leq r < n)$ will take $n-1$ “RD” steps (as shown in Figure 7) to read the $n-1$ index files generated by other processes. At each “RD” step $s(0 \leq s < n-1)$, process r reads index file $f = (n+r-s-1)\%n$ and loads the index record into its own in-memory index. After s “RD” steps, each process will have a full copy of the entire metadata index. As shown in Figure 7, at RD step 0 and 1, process 0 reads index file 2 and 1 respectively, process 1 reads index file 0 and 2 respectively, and process 2 reads index file 1 and 0 respectively. In this way, we can guarantee that no index files will be accessed by multiple processes simultaneously. After the retrospective loading procedure, each process will have a full copy of the metadata index on the entire self-describing file collection.

3.4.2 *Index Reloading.* MIQS stores the metadata index in the MIQS index file. Therefore, when applications need to be restarted or deployed at another HPC facility, MIQS can simply reload the index files and get ready for query processing. By following the data

layout in the MIQS index file format, the MIQS index builder reads each attribute frame, retrieves the corresponding index records and loads them into the in-memory index. Utilizing the same index file loading procedure as in retrospective loading, MIQS performs one more RD step to load a full copy of the metadata index for each process when performing index reloading. For example, as shown in Figure 7, three processes have stored their in-memory index into three different index files at the WR step when writing in-memory index to index files. After shifting the index file read in the consequent two RD steps by following the same procedure as in retrospective loading, process 0 loads file 0 at RD step 2, process 1 loads file 1 and process 2 loads file 2. Once completed, all three processes have the same full copy of the metadata in-memory index.

3.5 Serving Queries

In order to process metadata queries, MIQS follows the procedure described in Algorithm 2. MIQS first searches the root level of the ART to locate leaf node $leaf_k$ which has a matching attribute name. Then, based on the data type of the attribute value, MIQS starts from leaf node $leaf_k$ and searches through either the ART or the SBST associated with $leaf_k$. When searching the appropriate tree for attribute values, the relationship R is passed to the search function in order to find appropriate leaf nodes which satisfy the specified relationship between itself and the given attribute value. Once the leaf nodes are found, the lists of tuples containing the path identifiers to the owner files and the owner object can be retrieved. After looking up the file path list and the object path list with the identifier of each file path or object path, the actual file paths and object paths are retrieved and returned as the search result.

Algorithm 2 Serving Query

```

1: procedure SERVINGQUERY(key, R, value, index_anchor)
2:    $ART_{root} \leftarrow index\_anchor$ 
3:    $leaf_k \leftarrow search\_ART(key, ART_{root})$ 
4:   if value is a string then
5:      $ART_{value} \leftarrow get\_ART(leaf_k)$ .
6:      $leaf_v \leftarrow search\_ART(value, R, ART_{value})$ .
7:      $result \leftarrow get\_tuple\_list(leaf_v)$ .
8:   else if value is a number then
9:      $SBST_{value} \leftarrow get\_SBST(leaf_k)$ .
10:     $leaf_v \leftarrow search\_SBST(value, R, SBST_{value})$ 
11:     $result \leftarrow get\_tuple\_list(leaf_v)$ .
12:   end if
13:   return result
14: end procedure

```

Each process with MIQS can perform metadata searching, and, as a result of the shared-nothing architecture of MIQS, no communication is needed among these processes. Consequently, the throughput of the search performance is expected to be proportional to the number of processes, which makes the metadata search process highly scalable. Also, as there is no communication with other data sources like DBMS, the search latency is expected to be small.

4 EVALUATION

MIQS is designed to provide an efficient metadata search service that complies with the self-contained data management paradigm. In this section, we show how MIQS is suitable for the self-describing and self-contained data management solution by reporting the evaluation results of index construction performance, query performance, and storage overhead. In our evaluation, we compare MIQS with a MongoDB-powered metadata search solution which we consider to be a typical example of DBMS-based solution as the document-based data model of the MongoDB is more flexible than RDBMS-based solutions and hence is better suited for both hierarchical and flat metadata organization. As such, we do not compare against RDBMS solutions in this paper. Also, for fair comparison, we evaluate both solutions by building metadata index while scanning a set of HDF5 files, since we did not find any in-situ index construction use case through our investigation study on MongoDB solutions.

4.1 Experimental Setup

4.1.1 Platform. We conducted our evaluation of MIQS on the Edison supercomputer hosted at the National Energy Research Scientific Computing Center (NERSC). This system consists of 5586 “Ivy Bridge” compute nodes, where each node features two 12-core Intel® “Ivy Bridge” processors at 2.4GHz and 64 GB of DDR3 1866 memory. The Edison employs a Cray Aries interconnect with Dragonfly topology and 23.7 TB/s global bandwidth. The compute nodes use GPFS for its home directory and multiple Lustre file systems as scratch spaces. We used a 30 PB Lustre file system with over 700 GB/s peak I/O bandwidth for our evaluation.

For the comparison experiment against MongoDB, we used the MongoDB instance that is also maintained by NERSC. It is installed on a separate machine that can be accessed from the Edison compute nodes through an Infiniband network with 56 Gb/s bandwidth. The host machine of MongoDB has two 16-core Intel® Xeon™ processors E5-2698 v3 (“Haswell”) at 2.3 GHz and 128 GB of DDR4 2133 MHz memory. It features a 6 TB 7200rpm HDD with 6 Gb/s SAS interface. The MongoDB was initialized with the default configuration, which enables data compression on top of the WiredTiger storage engine [44].

4.1.2 Dataset. Our dataset contains a set of 100 real-world HDF5 files coming from the Baryon Oscillation Spectroscopic Survey (BOSS) [5]. The total data size of these 100 HDF5 files on disk is approximately 145GB with 144 million attribute key-value pairs attached to 1.5 million data objects (roughly 96 attribute key-value pairs on each object). Figure 8 summarizes the statistics of this dataset. The sizes of the files range from 400MB to 2.2GB (shown in Figure 8a). The number of objects per file spans between 4,800 to 23,000 (shown in Figure 8b). With over 250 different attributes attached to different data objects in each file, the number of attribute key-value pairs per file ranges from 300,000 to 2,300,000 (Figure 8c).

4.1.3 Evaluation Procedure. In our evaluation, we built two test drivers - the MongoDB test driver and the MIQS test driver. Both test drivers are MPI programs implemented with metadata indexing and querying functions. For the evaluation on MongoDB and MIQS,

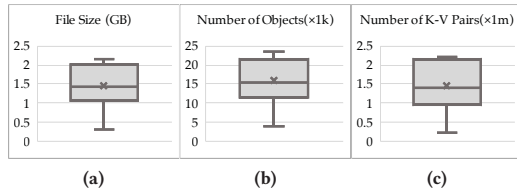


Figure 8: Statistics of datasets

we performed 5 test rounds and for each test round $r = \{1, 2, 3, 4, 5\}$ we ran $20 \times r$ processes for each test driver and we bound each process to a single core of the CPU and a different HDF5 file. As there are two 12-core CPUs on each compute node, each test round r uses r compute nodes for running $20 \times r$ processes.

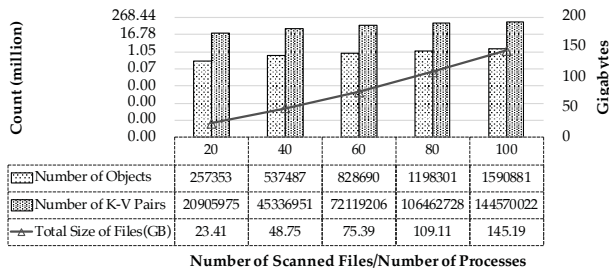


Figure 9: Dataset size at each test round

Figure 9 shows the size of the dataset being scanned at each test round, including the number of scanned objects and the number of scanned attribute key-value pairs in log-scale (the left log-scale axis with base 2), and the total size of the indexed files in gigabytes (the right axis).

For the evaluation on MongoDB, due to the schema constraints discussed in Section 2.2, it is impossible to know the paths to the owner objects of all the attributes without scanning the entire dataset. Therefore, we selected 16 representative metadata attributes with known owner object paths for creating attribute indexes. As shown in the first and second column of Table 2, these 16 attributes include 5 most frequently occurring integer attributes, 5 most frequently occurring float attributes, and 6 most frequently occurring string attributes. We consider the number and the diversity of these attributes to be sufficient for simulating real-world metadata search scenarios.

In order to build the metadata index on MongoDB at each test round, we first sent index creation commands to MongoDB to create indexes on these 16 attributes. Each test driver process then scanned a single HDF5 file and transformed the HDF5 metadata into several BSON documents that are smaller than 16MB [1]. These documents were then inserted into MongoDB. We believe, based on our study of MongoDB-based solutions, that this practice is necessary and in accordance with the realistic MongoDB-based use cases.

We evaluated MIQS for 10 test rounds. In the first 5 rounds, we ran MIQS test drivers with the settings shown in Figure 9, and created indexes for the 16 attributes we selected in Table 2. For the second set of five test rounds, we used the same settings but we let

Table 2: Attributes and sample query conditions

| Data Type | Attributes | Sample Query Conditions |
|----------------|------------|--|
| Integer | BESTEXP | BESTEXP=103179 |
| | DARKTIME | DARKTIME=0 |
| | BADPIXEL | BADPIXEL= 155701 |
| | COLLB | COLLB=26660 |
| | HIGHREJ | HIGHREJ=8 |
| Floating-point | HELIO_RV | HELIO_RV=26.6203 |
| | IOFFSTD | IOFFSTD=0.0133138 |
| | CRVAL1 | CRVAL1=3.5528 |
| | M1PISTON | M1PISTON=661.53 |
| | FBADPIX2 | FBADPIX2=0.231077 |
| String | AUTHOR | AUTHOR="Scott Burles & David Schlegel" |
| | FILENAME | FILENAME="badpixels-56149-b1.fits.gz" |
| | EXPOSURE | EXPOSURE="sdR-b2-00154990.fit" |
| | LAMPLIST | LAMPLIST="lampghcdne.dat" |
| | COMMENT | COMMENT="sp2blue cards follow" |
| | DAQVER | DAQVER="1.2.7" |

the test driver index every attribute it encountered while scanning the HDF5 files.

Throughout our evaluation, we refer to the index construction procedure as “indexing”, a file or attribute that has been scanned as “scanned file” or “scanned attribute”, and a file or attribute that has been successfully indexed as “indexed file” or “indexed attribute”.

After indexing, both test drivers performed 1,024 metadata queries from each process against their target data sources - MongoDB and MIQS indexes, respectively. Each query asked for identifiers of objects that match with the given query condition. Our query conditions covered the 16 selected attributes with non-repetitive attribute values of three different data types - integer, floating-point, and string. Thus, these 1,024 metadata queries covered 64 different values for each selected attribute. The last column in Table 2 shows selected sample query conditions for each data type from the 1,024 queries we issued.

4.2 Index Construction Time

Index construction time is very important to a metadata search service for large-scale self-describing datasets. It determines the time for getting metadata index ready to use and hence has a significant impact on the entire software life cycle when the metadata search service is integrated.

Figure 10 shows the index construction time for MongoDB and MIQS during the 5 test rounds. For MongoDB, the total indexing time needed to create indexes for the 16 specified attributes ranged from 373 to 587 seconds. In comparison, when indexes were constructed on the same set of 16 attributes in MIQS, the indexing time ranged from 194 seconds to 310 seconds, which is approximately half of the indexing time of MongoDB. We also tested indexing all attributes in MIQS and the indexing time ranged from 533 seconds to 860 seconds. Finally, we tested building the in-memory index from the MIQS index files. The time required for loading 16-attribute index files varied from 3 seconds to 27 seconds, approximately a 95%-99% time reduction compared to MongoDB-powered solution. Moreover, MIQS was able to load the index files for all attributes in

10.3 - 112 seconds, which is equivalent to only 2% - 19% of the time taken by the MongoDB-powered solution.

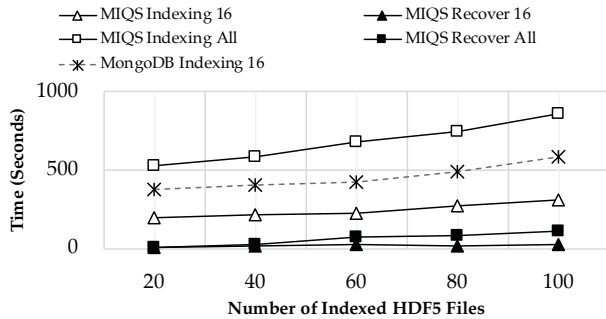


Figure 10: Time for index construction

For a better understanding on this result, we further decompose the indexing time. As shown in Figure 11, the time spent scanning HDF5 files was roughly the same, approximately 200 seconds, for both the MongoDB solution and MIQS, regardless of the number of indexed attributes (Figure 11a, 11b, and 11c). However, the MongoDB test driver also spent 200 to 359 seconds for inserting the metadata BSON [7] documents into MongoDB via network. The total time spent inserting BSON documents is 1.5 - 2x the time spent scanning HDF5 files and is significant given only 16 attributes were indexed (Figure 11a). In fact, this time may include time spent transferring BSON documents over the network and performing I/O operations at the MongoDB server.

The time for inserting BSON documents was even 1.5 to 2 times of the time spent for scanning the HDF5 data files, and this may include the time for transferring the BSON document over the network and also the I/O operation occurred at MongoDB server. Apart from the data file scanning procedure, building the actual in-memory indexes for 16 attributes only took 24-32 seconds for MIQS, while index persistence and index retrospective loading only took 3 to 28 seconds (Figure 11b). For indexing all attributes, MIQS spent 287 to 470 seconds building the in-memory indexes and the index persistence and retrospective loading only took 21 to 114 seconds (Figure 11c). This is quite comparable to the metadata indexing time in MongoDB; however, MIQS built indexes for all attributes while only 16 attributes were indexed in MongoDB.

Figure 11d illustrates the case where self-describing data files are transferred to a new site and the in-memory index has to be reloaded from the index files in MIQS. It took less than 40 seconds to load metadata indexes on 16 attributes utilizing up to 100 processes to initialize the metadata search service. The metadata search service can be initialized within 2 minutes when loading the metadata indexes of all attributes.

Overall, our evaluation on index construction time indicates that MIQS requires much less time than the MongoDB-powered solutions. It can also be rapidly deployed in the case where metadata index files already exist. Consequently, applications with MIQS can begin to perform metadata query much earlier than with MongoDB.

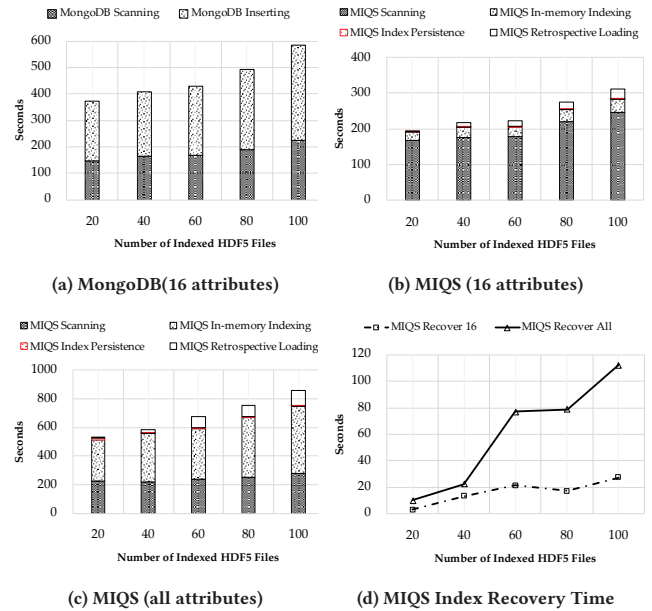


Figure 11: Breakdown analysis of indexing time

4.3 Query Performance

Query performance is critical to a metadata search service. By utilizing an in-memory shared-nothing parallel architecture, MIQS delivers an efficient metadata search service.

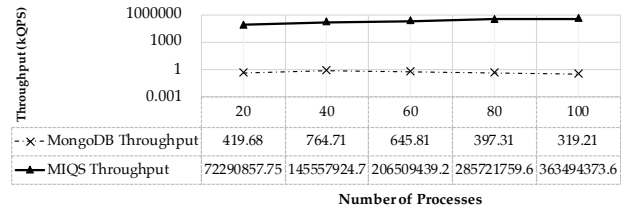


Figure 12: Query throughput comparison (kQPS)

As reported in Figure 12, for an application with 20 processes, the overall query throughput of MIQS was 72.2 billion queries per second while the MongoDB only achieved 419.7k queries per second. This is a drastic performance improvement of over 172,000x speedup. As shown in the figure, the query throughput of MIQS scales almost linearly with the number of processes, and ranges from 72.2 billion QPS all the way up to 363.4 billion QPS. However, the MongoDB query throughput only slightly increased to 764.7 kQPS when queries were issued by 40 processes and dropped from 645.8 kQPS to 319.2 kQPS when the number of processes increased from 60 to 100.

The query latency is depicted in Figure 13. Although the latency of each query in MongoDB remains steady around 0.28 to 0.29 milliseconds, as the number of processes increases from 20 to 100, the query latency dramatically increases from 47 seconds all the way up to 313 seconds. This is almost equal to the time

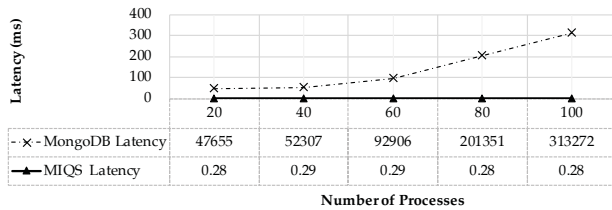


Figure 13: Query latency comparison(ms)

spent scanning 100 HDF5 data files during our index construction experiment. These results confirm that, as a benefit of following the self-contained data management paradigm, MIQS outperforms MongoDB-powered solutions and delivers superior query performance for metadata searches.

4.4 Storage Consumption

MIQS is designed as an in-memory metadata search service wherein processes have their own full copy of the metadata index. Thus, the memory footprint has to be small. In addition, the storage consumption must be minimized so that the index files can be integrated into existing self-describing file collections without significantly increasing the size of the dataset.

In our evaluation, we collected MongoDB storage consumption statistics using the administrative command “dbStats”. The result of this command includes the “indexSize” field which represents the total size of the indexes and also the “dataSize” field which represents the overall storage consumption of MongoDB. The storage overhead for maintaining BSON documents in MongoDB is then given by the difference between “dataSize” and “indexSize”. For the memory consumption of MIQS, we track every memory allocation required by the indexing data structures. For storage consumption, we measure and analyze the size of MIQS index files.

Figure 14 shows the storage consumption comparison between MongoDB and MIQS. As shown in Figure 14a, when the index is built against only 16 attributes, the total storage consumption of MongoDB ranges from about 680 MB to approximately 4.2 GB. Regardless of scale, the index size in MongoDB occupies only about 20% of the overall storage space consumption while 80% is used to store metadata BSON documents.

When it comes to MIQS, as depicted in Figure 14b, the total storage consumption for indexing 16 attributes does not exceed 0.6 GB at any scale. The index files consume roughly the same amount of space on disk as the in-memory index consumes in the memory. For our largest test, MIQS reduces the total storage consumption by about 75.4% when compared to MongoDB.

Indexing all attributes in MIQS yields a total storage consumption that ranges from 1.5GB to 7.8 GB where the in-memory index consumes half of the total storage space (Figure 14c). Although the total storage consumption of MIQS is twice that of MongoDB, MIQS indexes all attributes while MongoDB indexes only 16. Further, MongoDB wastes storage space storing duplicated metadata that already exists in the self-describing data files.

Figure 14d shows the overall storage consumption in parallel scenarios where the number of processes increases from 20 to 100.

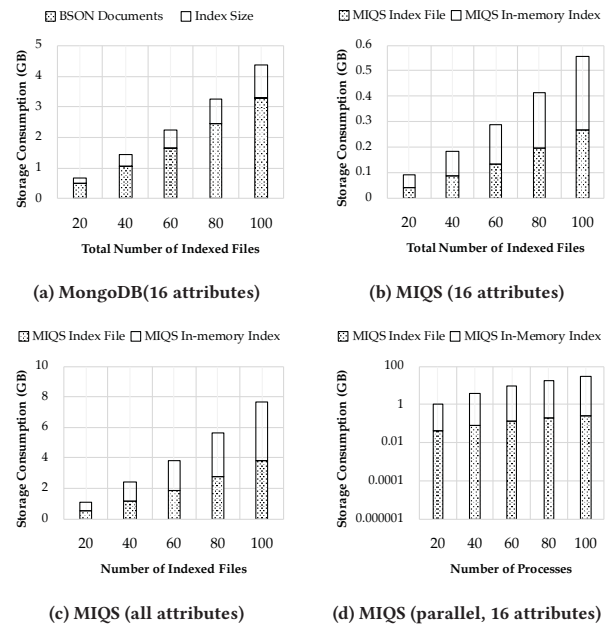


Figure 14: Breakdown analysis on memory consumption

As illustrated in the figure, parallel applications with MIQS consume up to approximately 30 GB of storage space when indexing 16 attributes. However, considering that each process only needs to maintain around 600 MB of memory for metadata search service, which provides the high query throughput we have reported in the previous section, we believe the storage consumed by the index is worthwhile. The low storage consumption of MIQS further reinforces the idea that the design of MIQS complies with self-contained data management paradigm.

4.5 Summary of Evaluation

MIQS is designed as a metadata indexing and querying service that follows the self-contained data management paradigm. From our evaluation, we can discern three significant benefits that result from adherence to this paradigm: schema-free metadata indexing (MIQS was capable of indexing and querying all attributes but MongoDB was not), rapid index construction (hundreds of seconds saved in MIQS as compared to MongoDB) and superior search performance (172k× throughput improvement compared to MongoDB-powered solution). Our evaluation also shows that MIQS takes only a small amount of memory and that the small size of its index files enhances portability. Our results demonstrate the promise of MIQS and present a rationale for integration with existing self-describing and self-contained data management solutions.

5 CONCLUSION & FUTURE WORK

Self-describing data formats store metadata alongside the data objects themselves. Existing metadata search services for these self-describing data formats are primarily built upon external database management systems (DBMS). Unfortunately, these solutions disregard the principle of self-describing and self-contained data

management paradigm, hinder schema-transparency and portability and hence introduce additional complexity in deployment and maintenance.

In this paper, we have proposed the Metadata Indexing and Querying Service (MIQS) for efficient metadata search on self-describing file formats such as HDF5. MIQS provides portable and schema-free indexing and querying solution that complies with the self-contained data management paradigm. It eliminates the necessity of maintaining a disjoint data source that duplicates metadata and thereby requires a smaller overall storage footprint, as compared to the DBMS-powered solutions. We have conducted extensive evaluations to compare MIQS against a state-of-the-art DBMS-based metadata querying solution - the MongoDB-powered metadata querying service. MIQS significantly outperforms MongoDB-powered solution in both index construction and metadata search. More importantly, our evaluation shows that MIQS complies with the self-contained data management principle of self-describing data formats.

At the time of writing this paper, we are in contact with the HDF5 developers to make MIQS available as an indexing and querying component for the HDF5 library. In the future, we will enhance MIQS to support complex queries on compound data types in addition to numeric values or strings.

ACKNOWLEDGMENTS

We are thankful to the anonymous reviewers for their valuable feedback. This research is supported in part by the National Science Foundation under grant CNS-1338078, CNS-1362134, CCF-1409946, CCF-1718336, OAC-1835892, and CNS-1817094. This work is supported in part by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. (Project: EOD-HDF5: Experimental and Observational Data enhancements to HDF5, Program managers: Dr. Laura Biven and Dr. Lucy Nowell). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a DOE Office of Science User Facility.

REFERENCES

- [1] 2018. MongoDB Limits and Thresholds. <https://docs.mongodb.com/manual/reference/limits/>.
- [2] 2018. Simple and fast C library implementing a thread-safe API to manage hash-tables, linked lists, lock-free ring buffers and queues. <https://github.com/xant/libhl>.
- [3] MG Aartsen, K Abraham, M Ackermann, J Adams, JA Aguilar, M Ahlers, M Ahrens, D Altmann, K Andeen, T Anderson, et al. 2016. Search for sources of High-Energy neutrons with four years of data from the IceTop Detector. *The Astrophysical Journal* 830, 2 (2016), 129.
- [4] MG Aartsen, M Ackermann, J Adams, JA Aguilar, Markus Ahlers, M Ahrens, I Al Samarai, D Altmann, K Andeen, T Anderson, et al. 2017. Constraints on galactic neutrino emission with seven years of IceCube data. *The Astrophysical Journal* 849, 1 (2017), 67.
- [5] Christopher P Ahn, Rachael Alexandroff, Carlos Allende Prieto, Scott F Anderson, Timothy Anderton, Brett H Andrews, Éric Aubourg, Stephen Bailey, Eduardo Balbinot, Rory Barnes, et al. 2012. The ninth data release of the Sloan Digital Sky Survey: first spectroscopic data from the SDSS-III Baryon Oscillation Spectroscopic Survey. *The Astrophysical Journal Supplement Series* 203, 2 (2012), 21.
- [6] Shadab Alam, Metin Ata, Stephen Bailey, Florian Beutler, Dmitry Bizyaev, Jonathan A Blazek, Adam S Bolton, Joel R Brownstein, Angela Burden, Chia-Hsun Chuang, et al. 2017. The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample. *Monthly Notices of the Royal Astronomical Society* 470, 3 (2017), 2617–2652.
- [7] bsonspec.org. 2018. BinAjary JSON Specification. <http://bsonspec.org/spec.html>.
- [8] Chi Chen, Zhi Deng, Richard Tran, Hanmei Tang, Iek-Heng Chu, and Shyue Ping Ong. 2017. Accurate force field for molybdenum by machine learning large materials data. *Physical Review Materials* 1, 4 (2017), 043603.
- [9] Tull Craig E., Essiari Abdelilah, Gunter Dan, et al. 2013. The SPOT Suite project. <http://spot.nersc.gov/>.
- [10] Digital Curation Conference (DCC). 2018. Scientific Metadata. <http://www.dcc.ac.uk/resources/curation-reference-manual/chapters-production/scientific-metadata>.
- [11] Jeffrey J Donatelli, James A Sethian, and Peter H Zwart. 2017. Reconstruction from limited single-particle diffraction data via simultaneous determination of state, orientation, intensity, and phase. *Proceedings of the National Academy of Sciences* 114, 28 (2017), 7222–7227.
- [12] Bin Dong, Surendra Byna, and Kesheng Wu. 2015. Spatially clustered join on heterogeneous scientific data sets. In *2015 IEEE International Conference on Big Data (Big Data)*. IEEE, 371–380.
- [13] Mike Folk, Albert Cheng, and Kim Yates. 1999. HDF5: A file format and I/O library for high performance computing applications. In *Proceedings of supercomputing*, Vol. 99, 5–33.
- [14] P Greenfield, M Droettboom, and E Bray. 2015. ASDF: A new data format for astronomy. *Astronomy and Computing* 12 (2015), 240–251.
- [15] The HDF Group. 2018. HDF5 Topic Parallel Indexing Branch. <https://git.hdfgroup.org/users/jsoumagne/repos/hdf5/browse?at=refs%2Fheads%2Ftopic-parallel-indexing>.
- [16] The HDF Group. 2018. HDF5 Users. <https://support.hdfgroup.org/HDF5/users5.html>.
- [17] Kohei Hiraga, Osamu Tatebe, and Hideyuki Kawashima. 2018. PPMDS: A Distributed Metadata Server Based on Nonblocking Transactions. In *Fifth International Conference on Social Networks Analysis, Management and Security, SNAMS 2018, Valencia, Spain, October 15-18, 2018*. 202–208. <https://doi.org/10.1109/SNAMS.2018.8554478>
- [18] Joint Genome Institute. 2013. The JGI Archive and Metadata Organizer (JAMO). <http://cs.lbl.gov/news-media/news/2013/new-metadata-organizer-streamlines-jgi-data-management>.
- [19] json.org. 2018. Introducing JSON. <https://www.json.org>.
- [20] Daniel Korenblum, Daniel Rubin, Sandy Napel, Cesar Rodriguez, and Chris Beaulieu. 2011. Managing biomedical image metadata for search and retrieval of similar images. *Journal of digital imaging* 24, 4 (2011), 739–748.
- [21] Margaret Lawson and Jay Lofstead. 2018. Using a Robust Metadata Management System to Accelerate Scientific Discovery at Extreme Scales. In *Proceedings of the 2nd PDSW-DISCS '18*. <https://doi.org/10.1109/PDSW-DISCS.2018.00005>
- [22] Viktor Leis, Alfons Kemper, and Thomas Neumann. 2013. The Adaptive Radix Tree: ARTful Indexing for Main-memory Databases. In *Proceedings of the 2013 IEEE International Conference on Data Engineering (ICDE 2013) (ICDE '13)*. IEEE Computer Society, Washington, DC, USA, 38–49. <https://doi.org/10.1109/ICDE.2013.6544812>
- [23] Andrew W Leung, Minglong Shao, Timothy Bisson, Shankar Pasupathy, and Ethan L Miller. 2009. Spyglass: Fast, Scalable Metadata Search for Large-Scale Storage Systems. In *FAST*, Vol. 9, 153–166.
- [24] Jialin Liu, Debbie Bard, Quincey Koziol, Stephen Bailey, et al. 2017. Searching for millions of objects in the BOSS spectroscopic survey data with H5Boss. In *2017 New York Scientific Data Summit (NYSDS)*. 1–9. <https://doi.org/10.1109/NYSDS.2017.8085044>
- [25] Yaning Liu, George Shu Heng Pau, and Stefan Finsterle. 2017. Implicit sampling combined with reduced order modeling for the inversion of vadose zone hydrological data. *Computers & Geosciences* (2017).
- [26] Jay F. Lofstead, Scott Klasky, et al. 2008. Flexible IO and integration for scientific codes through the adaptable IO system (ADIOS). In *CLADE*. 15–24.
- [27] Arun Mannodi-Kanakkithodi, Tran Doan Huan, and Rampi Ramprasad. 2017. Mining materials design rules from data: The example of polymer dielectrics. *Chemistry of Materials* 29, 21 (2017), 9001–9010.
- [28] Vilobh Meshram, Xavier Besseron, Xiangyong Ouyang, Raghunath Rajachandrasekar, Ravi Prakash, and Dhableswar K. Panda. 2011. Can a Decentralized Metadata Service Layer Benefit Parallel Filesystems?. In *2011 IEEE International Conference on Cluster Computing (CLUSTER)*, Austin, TX, USA, September 26-30, 2011. 484–493. <https://doi.org/10.1109/CLUSTER.2011.85>
- [29] MongoDB. 2018. MongoDB. <https://www.mongodb.com>.
- [30] mongodb.com. 2018. The MongoDB 4.0 Manual. <https://docs.mongodb.com/manual/>.
- [31] David Paez-Espino, I Chen, A Min, Krishna Palaniappan, Anna Ratner, Ken Chu, Ernest Szeto, Manoj Pillay, Jinghua Huang, Victor M Markowitz, et al. 2017. IMG/VR: a database of cultured and uncultured DNA Viruses and retroviruses. *Nucleic acids research* 45, D1 (2017), D457–D465.
- [32] PostgreSQL. 2018. PostgreSQL. <https://www.postgresql.org>.
- [33] Russ Rew and Glenn Davis. 1990. NetCDF: an interface for scientific data access. *IEEE computer graphics and applications* 10, 4 (1990), 76–82.
- [34] Frank B Schmuck and Roger L Haskin. 2002. GPFS: A Shared-Disk File System for Large Computing Clusters. In *FAST*, Vol. 2.

- [35] Philip Schwan et al. 2003. Lustre: Building a file system for 1000-node clusters. In *Proceedings of the 2003 Linux symposium*, Vol. 2003. 380–386.
- [36] Self-balancing binary search tree. 2019. Self-balancing binary search tree – Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/wiki/Self-balancing_binary_search_tree [Online; accessed 10-April-2019].
- [37] Arie Shoshani and Doron Rotem. 2009. *Scientific data management: challenges, technology, and deployment*. Chapman and Hall/CRC.
- [38] Hyogi Sim, Youngjae Kim, Sudharshan S. Vazhkudai, Geoffroy R. Vallée, Seung-Hwan Lim, and Ali Raza Butt. 2017. Tagit: an integrated indexing and search service for file systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2017, Denver, CO, USA, November 12 - 17, 2017*, 5:1–5:12. <https://doi.org/10.1145/3126908.3126929>
- [39] sqlite.org. 2017. SQLite. <https://sqlite.org>.
- [40] Houjun Tang, Suren Byna, Bin Dong, Jialin Liu, and Quincey Koziol. 2017. SoMeta: Scalable Object-Centric Metadata Management for High Performance Computing. In *Cluster Computing (CLUSTER), 2017 IEEE International Conference on*. IEEE, 359–369.
- [41] Alexander Thomson and Daniel J. Abadi. 2015. CalvinFS: Consistent WAN Replication and Scalable Metadata Management for Distributed File Systems. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies, FAST 2015, Santa Clara, CA, USA, February 16-19, 2015*, Jiri Schindler and Erez Zadok (Eds.). USENIX Association, 1–14. <https://www.usenix.org/conference/fast15/technical-sessions/presentation/thomson>
- [42] Teng Wang, Adam Moody, Yue Zhu, Kathryn Mohror, Kento Sato, Tanzima Islam, and Weikuan Yu. 2017. MetaKV: A Key-Value Store for Metadata Management of Distributed Burst Buffers. In *2017 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2017, Orlando, FL, USA, May 29 - June 2, 2017*. IEEE Computer Society, 1174–1183. <https://doi.org/10.1109/IPDPS.2017.39>
- [43] Zeyi Wen, Xingyang Liu, Hongjian Cao, and Bingsheng He. 2018. RTSI: An Index Structure for Multi-Modal Real-Time Search on Live Audio Streaming Services. In *34th IEEE International Conference on Data Engineering, ICDE 2018, Paris, France, April 16-19, 2018*. 1495–1506. <https://doi.org/10.1109/ICDE.2018.00168>
- [44] WiredTiger. 2018. WiredTiger. <http://www.wiredtiger.com/>.
- [45] Quanqing Xu, Rajesh Vellore Arumugam, Khai Leong Yang, and Sridhar Mahadevan. 2013. Drop: Facilitating distributed metadata management in eb-scale storage systems. In *2013 IEEE 29th symposium on mass storage systems and technologies (MSST)*. IEEE, 1–10.
- [46] Wei Zhang, Houjun Tang, Suren Byna, and Yong Chen. 2018. DART: Distributed Adaptive Radix Tree for Efficient Affix-based Keyword Search on HPC Systems. In *Proceedings of The 27th International Conference on Parallel Architectures and Compilation Techniques (PACT'18)*. <https://doi.org/10.1145/3243176.3243207>
- [47] Dongfang Zhao, Kan Qiao, Zhou Zhou, Tonglin Li, Zhihan Lu, and Xiaohua Xu. 2017. Toward Efficient and Flexible Metadata Indexing of Big Data Systems. *IEEE Trans. Big Data* 3, 1 (2017), 107–117.
- [48] Qing Zheng, Charles D. Cranor, Danhao Guo, Gregory R. Ganger, George Amvrosiadis, Garth A. Gibson, Bradley W. Settlemyer, Gary Grider, and Fan Guo. 2018. Scaling embedded in-situ indexing with deltaFS. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2018, Dallas, TX, USA, November 11-16, 2018*. 3:1–3:15. <http://dl.acm.org/citation.cfm?id=3291660>

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

We ran our experiment on NERSC's Edison supercomputer and our MongoDB instance is also provided by NERSC database services. Both of our MongoDB test driver and MIQS test driver are written in C and are deployed on one-process-per-computing-node basis. We compare the indexing performance, search performance and storage consumption of MongoDB and MIQS in our experiment, and the result is reported as written in the paper.

ARTIFACT AVAILABILITY

Software Artifact Availability: Some author-created software artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Hardware Artifact Availability: Some author-created hardware artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Data Artifact Availability: Some author-created data artifacts are NOT maintained in a public repository or are NOT available under an OSI-approved license.

Proprietary Artifacts: No author-created artifacts are proprietary.

List of URLs and/or DOIs where artifacts are available:

<https://github.com/zhangwei217245/HDF5Meta>

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Cray XC30 supercomputer/2*12-core Intel "Ivy Bridge" processor at 2.4 GHz/Each node has 64 GB DDR3 1866 MHz memory (four 8 GB DIMMs per socket)/Lustre file system peak @ 700GB/s

Operating systems and versions: SuSE Linux Enterprise Server running lightweight Linux Kernel 4.4.162

Compilers and versions: GCC 7.3.0

Libraries and versions: HDF5 1.10.2.

Key algorithms: Adaptive Radix Tree and Red-Black Tree and their corresponding operations

Input datasets and versions: HDF5 files generated by Baryon Oscillation Spectroscopic Survey

Output from scripts that gathers execution environment information.

```
SLURM_NODELIST=nid0[1572-1573]
LESSKEY=/etc/lesskey.bin
PE_LIBSCI_VOLATILE_PRGENV=CRAY GNU INTEL
PE_SMA_DEFAULT_PKGCONFIG_VARIABLES=PE_SMA_COMPFLAG_@_
↪ prgenv@
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_mic_knl=160
MODULE_VERSION_STACK=3.2.10.6
```

```
GREP_COLOR=1;32
KSH_AUTOLOAD=1
ZAP_LIBPATH=/opt/ovis/lib64/ovis-lib
SLURM_CHECKPOINT_IMAGE_DIR=/var/slurm/checkpoint
SLURM_JOB_NAME=sh
GNU_VERSION=7.3.0
MANPATH=/usr/common/software/man:/usr/common/mss/man_
↪ :/usr/common/nsg/man:/opt/java/jdk1.8.0_51/man:/_
↪ opt/cray/pe/atp/2.1.3/man:/opt/cray/alps/6.6.43-_
↪ 6.0.7.0_26.4__ga796da3.ari/man:/opt/cray/job/2.2_
↪ .3-6.0.7.0_44.1__g6c4e934.ari/man:/opt/cray/pe/p_
↪ mi/5.0.14/man:/opt/cray/pe/libsci/18.07.1/man:/o_
↪ pt/cray/pe/man/csm1version:/opt/cray/pe/craype/2_
↪ .5.15/man:/opt/gcc/7.3.0/snos/share/man:/opt/cra_
↪ y/pe/mpt/7.7.3/gni/man/mpich:/usr/syscom/nsg/man_
↪ :/opt/cray/pe/modules/3.2.10.6/share/man:/usr/lo_
↪ cal/man:/usr/share/man:/opt/cray/share/man:/opt/_
↪ cray/share/man:/opt/cray/pe/man:/opt/cray/share/_
↪ man
PE_CXX_PKGCONFIG_LIBS=mpichcxx
PE_PETSC_DEFAULT_GENCOMPS_CRAY_skylake=86
PE_PAPI_DEFAULT_ACCEL_FAMILY_LIBS_nvidia=-lcupti,-l_
↪ cudart,-lcuda
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
NNTPSERVER=news
DVS_MAXNODES=1__
COLOR_NC=\e[0m
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_HDF5_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
XDG_SESSION_ID=69545
PE_TRILINOS_DEFAULT_GENCOMPS_CRAY_x86_64=87
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_FFTW_DEFAULT_TARGET_mic_knl=mic_knl
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
HDF5_DIR=/opt/cray/pe/hdf5-parallel/1.10.2.0/GNU/7.1
HOSTNAME=edison10
PE_HDF5_PARALLEL_VOLATILE_PRGENV=GNU
CRAY_UDREG_INCLUDE_OPTS=-I/opt/cray/udreg/2.3.2-6.0._
↪ 7.0_33.18__g5196236.ari/include
SLURM_TOPOLOGY_ADDR=s15.s10.nid01572
SLURMD_NODENAME=nid01572
CRAY_SITE_LIST_DIR=/etc/opt/cray/pe/modules
PE_SMA_DEFAULT_COMPFLAG_GNU=-fcray-pointer
```

```

LIBRARYMODULES=acml:alps:cray-dwarf:cray-fftw:cray-g
↪ a:cray-hdf5:cray-hdf5-parallel:cray-libsci:cray-
↪ libsci_acc:cray-mpich:cray-mpich-abi:cray-mpich2
↪ :cray-netcdf:cray-netcdf-hdf5parallel:cray-para
↪ lel-netcdf:cray-petsc:cray-petsc-complex:cray-sh
↪ mem:cray-tpsl:cray-trilinos:cudatoolkit:fftw:ga:
↪ hdf5:hdf5-parallel:iobuf:libfast:netcdf:netcdf-h
↪ df5parallel:ntk:onesided:papi:petsc:petsc-comple
↪ x:pmi:tpsl:trilinos:xt-libsci:xt-mpich2:xt-mpt:x
↪ t-papi
XKEYSYMDB=/usr/X11R6/lib/X11/XKeysymDB
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/
↪ opt/cray/pe/parallel-netcdf/1.8.1.3/@PRGENV@/PE
↪ _PARALLEL_NETCDF_DEFAULT_GENCOMPS@/lib/pkgconfig
PE_NETCDF_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/
↪ pe/netcdf/4.6.1.3/@PRGENV@/PE_NETCDF_DEFAULT_GE
↪ NCOMPS@/lib/pkgconfig
PE_TRILINOS_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cra
↪ y/pe/trilinos/12.12.1.1/@PRGENV@/PE_TRILINOS_DE
↪ FAULT_GENCOMPS@/PE_TRILINOS_DEFAULT_TARGET@/lib
↪ /pkgconfig
SLURM_PRIO_PROCESS=0
RCLOCAL_BASEOPTS=true
PE_MPICH_ALTERNATE_LIBS_dpm=_dpm
PE_ENV=GNU
INTEL_LICENSE_FILE=28518@crayintel.licenses.nersc.go
↪ v:28518@intel.licenses.nersc.gov
PE_SMA_DEFAULT_COMPFLAG=
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_HDF5_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
COLOR_RED=\e[0;31m
COLOR_CYAN=\e[0;36m
SLURM_SRUN_COMM_PORT=62382
COLOR_PURPLE=\e[0;35m
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
SHELL=/bin/bash
PKGCONFIG_ENABLED=1
TERM=xterm-color
HOST=edison10
SLURM_JOB_QOS=premium
PE_HDF5_PARALLEL_VOLATILE_PKGCONFIG_PATH=/opt/cray/p
↪ e/hdf5-parallel/1.10.2.0/@PRGENV@/PE_HDF5_PARAL
↪ LEL_GENCOMPS@/lib/pkgconfig
PROFILEREAD=true
PE_PETSC_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_TPSSL_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2 7.1
↪ 6.1
CLICOLOR=1
HISTSIZE=1000
SLURM_PTY_WIN_ROW=36
PE_TRILINOS_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
PE_PETSC_DEFAULT_GENCOMPS_GNU_haswell=71 53 49
PE_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
PE_TPSSL_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53 49
TMPDIR=/tmp
COLOR_LIGHT_RED=\e[1;31m

```

```

CRAYPE_DIR=/opt/cray/pe/craype/2.5.15
SSH_CLIENT=172.73.67.84 53311 22
PE_PARALLEL_NETCDF_DEFAULT_VOLATILE_PRGENV=GNU
CRAY_UGNI_POST_LINK_OPTS=-L/opt/cray/ugni/6.0.14.0-6
↪ .0.7.0_23.1__gea11d3d.ari/lib64
PE_TPSSL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSCI
CRAY_XPMEM_POST_LINK_OPTS=-L/opt/cray/xpmem/2.2.15-6
↪ .0.7.1_5.10__g7549d06.ari/lib64
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_x86_skylake=160
PE_PETSC_DEFAULT_GENCOMPS_INTEL_haswell=160
SLURM_TOPOLOGY_ADDR_PATTERN=switch.switch.node
PE_HDF5_DEFAULT_VOLATILE_PRGENV=GNU
CRAY_MPICH2_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-gnu
↪ /7.1
ALTD_SELECT_OFF_USERS=
PE_PETSC_DEFAULT_GENCOMPS_CRAY_interlagos=86
ALT_LINKER=/usr/common/software/altd/2.0/bin/ld
COLOR_LIGHT_PURPLE=\e[1;35m
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/op
↪ t/cray/pe/hdf5-parallel/1.10.2.0/@PRGENV@/PE_HD
↪ F5_PARALLEL_DEFAULT_GENCOMPS@/lib/pkgconfig
COLOR_GREEN=\e[0;32m
PE_FFTW_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /fftw/3.3.8.1/@PE_FFTW_DEFAULT_TARGET@/lib/pkgco
↪ nfig
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PKGCONFIG_PA
↪ TH=/opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/@PRG
↪ ENV@/PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS@/l
↪ ib/pkgconfig
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_mic_knl=86
ALTD_SELECT_ON=0
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_GA_DEFAULT_VOLATILE_PRGENV=GNU
PE_LIBSCI_DEFAULT_GENCOMPS_GNU_x86_64=71 61 51 49
PMI_CONTROL_PORT=63186
PE_MPICH_GENCOMPS_GNU=71 51 49
MORE=-sI
PE_TPSSL_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_PKGCONFIG_PRODUCTS=PE_HDF5_PARALLEL:PE_LIBSCI:PE_
↪ MPICH
PE_MPICH_DEFAULT_GENCOMPILERS_GNU=7.1 5.1 4.9
FPATH=/opt/cray/pe/modules/3.2.10.6/init/sh_funcs/n
↪ o_redirect:/opt/cray/pe/modules/3.2.10.6/init/sh
↪ _funcs/no_redirect:/opt/cray/pe/modules/3.2.10.6
↪ /init/sh_funcs/no_redirect
SLURM_CPU_BIND_VERBOSE=quiet
SLURM_SPANK_NERSC_ZONESORT_INTERVAL=0
COLOR_BLUE=\e[0;34m
PE_TPSSL_64_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray
↪ /pe/tpsl/18.06.1/@PRGENV@64/@PE_TPSSL_64_DEFAULT_
↪ GENCOMPS@/PE_TPSSL_64_DEFAULT_TARGET@/lib/pkgcon
↪ fig
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_haswell=86
ALTD_VERBOSE=0
PE_PETSC_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIBSC
↪ I:PE_HDF5_PARALLEL:PE_TPSSL

```


MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

```

PE_PAPI_DEFAULT_ACCEL_LIBS=nvidia35,-lcupti,-lcuda
↪ t,-lcuda
PE_CRAY_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/parallel-netcdf/1.8.1.3/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/hdf5-parallel/1.10.2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/hdf5/1.1.0.2.0/CRAY/8.6/lib/pkgconfig:/opt/cray/pe/ga/5.3.0.8/CRAY/8.6/lib/pkgconfig
PE_HDF5_PARALLEL_FIXED_PRGENV=CRAY INTEL
PE_TRILINOS_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.7
PE_PETSC_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
SSH_TTY=/dev/pts/50
PE_FORTRAN_PKGCONFIG_LIBS=hdf5hl_fortran_parallel:hdf5_fortran_parallel:mpichf90
PE_LIBSCI_DEFAULT_OMP_REQUIRES_openmp=mp
PE_SMA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.7.3/gni/sma@PE_SMA_DEFAULT_DIR_DEFAULT64@/lib64/pkgconfig
ALLINEA_QUEUE_DLL=/opt/cray/pe/mpt/7.7.3/gni/mpich-gnu/7.1/lib/libtvm_pich.so.3.0.1
SLURM_CPU_BIND_LIST=0xFFFFFFFFFFFF
PE_TRILINOS_DEFAULT_GENCOMPS_INTEL_x86_64=160
COLOR_YELLOW=\e[1;33m
CRAY_MPICH_BASEDIR=/opt/cray/pe/mpt/7.7.3/gni
ALPS_APP_ID=12930506
SLURM_NNODES=2
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2
↪ 7.1 6.1 5.3 4.9
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_haswell=160
JRE_HOME=/opt/java/jdk1.8.0_51/jre
PE_TPSSL_64_DEFAULT_GENCOMPS_CRAY_x86_skylake=86
PE_HDF5_PARALLEL_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1
↪ 5.3 4.9
USER=USER
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
LS_COLORS=
PE_TPSSL_64_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
LD_LIBRARY_PATH=/usr/common/software/darshan/3.1.4/1/ib:/opt/cray/job/2.2.3-6.0.7.0_44.1__g6c4e934.ar/i/lib64:/opt/gcc/7.3.0/snos/lib64
CRAY_HDF5_PARALLEL_VERSION=1.10.2.0
PE_TRILINOS_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.3
↪ 5.1 4.9
PE_FFTW_DEFAULT_TARGET_interlagos=interlagos
PE_LIBSCI_DEFAULT_VOLATILE_PRGENV=CRAY GNU INTEL
PE_TPSSL_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_TRILINOS_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_RCA_POST_LINK_OPTS=-L/opt/cray/rca/2.2.18-6.0.7.0_33.3_g2aa4f39.ari/lib64
↪ -lrca
PE_PETSC_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↪ GNU64 INTEL INTEL64

```

```

PE_TPSSL_64_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2
↪ 7.1 5.3 4.9
PE_HDF5_DIR=/opt/cray/pe/hdf5-parallel/1.10.2.0
PE_PKGCONFIG_LIBS=darshan-runtime:hdf5_hl_parallel:hdf5_parallel:AtpSigHandler:cray-rca:libsci_mpi:libsci:mpich
PE_MPICH_FIXED_PRGENV=INTEL
PE_TPSSL_64_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_LIBSCI_PKGCONFIG_VARIABLES=PE_LIBSCI_OMP_REQUIRES
↪ _@openmp@:PE_SCI_EXT_LIBPATH:PE_SCI_EXT_LIBNAME
CSCRATCH=/global/cscratch1/sd/USER
CRAY_IAA_INFO_FILE=/tmp/cray_iaa_info.12930506
PE_TPSSL_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
XNLSPATH=/usr/share/X11/nls
PE_PETSC_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_PETSC_DEFAULT_GENCOMPS_GNU_interlagos=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_GNU_sandybridge=71 53 49
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_sandybridge=160
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_skylake=16.0
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_mic_knl=160
PE_LIBSCI_DEFAULT_GENCOMPS_CRAY_x86_64=86
MPICH_ABORT_ON_ERROR=1
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_PETSC_DEFAULT_GENCOMPS_GNU_mic_knl=53
MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-gnu/7.1
PE_TPSSL_64_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2
↪ 7.1 5.3 4.9
PE_TPSSL_64_DEFAULT_GENCOMPS_INTEL_sandybridge=160
LDAPTLS_REQCERT=never
PE_PAPI_DEFAULT_PKGCONFIG_VARIABLES=PE_PAPI_ACCEL_LIBS
↪ BS@accelerator@
ALTD_ON=1
SLURM_STEP_NUM_NODES=1
SLURM_JOBID=12930506
SSH_AUTH_SOCKET=/tmp/ssh-uct6jVryWH/agent.49929
PE_MPICH_FORTRAN_PKGCONFIG_LIBS=mpichf90
MPICH_MPIIO_DVS_MAXNODES=14
PE_FFTW_DEFAULT_TARGET_sandybridge=sandybridge
PE_TPSSL_64_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
CRAY_PRGENVGNU=loaded
HOSTTYPE=x86_64
PE_NETCDF_HDF5PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_HDF5_PARALLEL
PE_HDF5_PARALLEL_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
ATP_POST_LINK_OPTS=-Wl,-L/opt/cray/pe/atp/2.1.3/libA
↪ pp/
PE_FFTW_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
SRUN_DEBUG=3
RCLOCAL_PRGENV=true
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_mic_knl=5.3
PE_LIBSCI_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
gcc_already_loaded=0
PE_TPSSL_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53 49

```

```

PE_HDF5_PARALLEL_FORTRAN_PKGCONFIG_LIBS=hdf5hl_fortr
↪ an_parallel:hdf5_fortran_parallel
GCC_VERSION=7.3.0
PE_PRODUCT_LIST=CRAY_RCA:CRAY_ALPS:DVS:CRAY_XPMEM:CR
↪ AY_DMAPP:CRAY_PMI:CRAY_UGNI:CRAY_UDREG:CRAY_LIBS
↪ CI:CRAYPE:CRAYPE_IVYBRIDGE:GNU:GCC
FROM_HEADER=
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_LIBSCI_GENCOMPS_INTEL_x86_64=160
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
CRAY_MPICH_ROOTDIR=/opt/cray/pe/mpt/7.7.3
PAGER=less
COLOR_BLACK=\e[0;30m
SLURM_LAUNCH_NODE_IPADDR=10.128.6.142
ALPS_LLI_STATUS_OFFSET=1
JAVA_PATH=/opt/java/jdk1.8.0_51
PE_MPICH_GENCOMPILERS_CRAY=8.6
PE_PETSC_DEFAULT_GENCOMPS_INTEL_skylake=160
PE_LIBSCI_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1 4.9
PE_MPICH_MODULE_NAME=cray-mpich
PE_PETSC_DEFAULT_GENCOMPS_GNU_skylake=61
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
CSHEDIT=emacs
SLURM_STEP_ID=0
ALPS_APP_PE=0
PE_MPICH_DEFAULT_VOLATILE_PRGENV=CRAY_GNU
PE_MPICH_TARGET_VAR_nvidia20=-lcudart
PE_CRAY_FIXED_PKGCONFIG_PATH=/opt/cray/pe/hdf5-para
↪ l1/1.10.2.0/CRAY/8.6/lib/pkgconfig
XDG_CONFIG_DIRS=/etc/xdg
PE_LIBSCI_GENCOMPS_CRAY_x86_64=86
PE_TPSL_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_LIBSCI_DEFAULT_GENCOMPILERS_CRAY_x86_64=8.6
PE_TPSL_64_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_LIB
↪ SCI
LSCOLORS=ExFxCxDxBxegedabagacad
PE_TPSL_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49
CRAY_DMAPP_INCLUDE_OPTS=-I/opt/cray/dmapp/7.1.1-6.0.
↪ 7.0_34.3__g5a674e0.ari/include
↪ -I/opt/cray/gni-headers/5.0.12.0-6.0.7.0_24.1__g
↪ 3b1768f.ari/include
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPS_GNU=51 49
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
LIBGL_DEBUG=quiet
PE_NETCDF_DEFAULT_GENCOMPS_GNU=
DVS_VERSION=0.9.0
MINICOM=-c on
PE_LIBSCI_PKGCONFIG_LIBS=libsci_mpi:libsci

```

```

USERMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:PrgE
↪ nv-pathscale:PrgEnv-pgi:acml:alps:apprentice:app
↪ rentice2:atp:blcr:cce:chapel:cray-cddb:cray-fftw
↪ :cray-ga:cray-hdf5:cray-hdf5-parallel:cray-lgdb:
↪ cray-libsci:cray-libsci_acc:cray-mpich:cray-mpic
↪ h-compat:cray-mpich2:cray-netcdf:cray-netcdf-hdf
↪ 5parallel:cray-parallel-netcdf:cray-petsc:cray-p
↪ etsc-complex:cray-shmem:cray-snp launcher:cray-tp
↪ sl:cray-trilinos:craypat:craype:craypkg-gen:cuda
↪ toolkit:ddt:fftw:ga:gcc:hdf5:hdf5-parallel:intel
↪ :iobuf:java:lgdb:libfast:libsci_acc:mpich1:netcd
↪ f:netcdf-hdf5parallel:netcdf-nofsync:netcdf-nofs
↪ ync-hdf5parallel:ntk:onesided:papi:parallel-netc
↪ df:pathscale:perf tools:perf tools-lite:petsc:pets
↪ c-complex:pgi:pmi:stat:totalview:tpsl:trilinos:x
↪ t-asyncpe:xt-craypat:xt-lgdb:xt-libsci:xt-mpich2
↪ :xt-mpt:xt-papi:xt-shmem:xt-totalview
CRAY_LIBSCI_DIR=/opt/cray/pe/libsci/18.07.1
CRAY_LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
SLURM_TASKS_PER_NODE=48(x2)
PE_HDF5_PARALLEL_PKGCONFIG_LIBS=hdf5_hl_parallel:hdf
↪ 5_parallel
PE_PKGCONFIG_DEFAULT_PRODUCTS=PE_TRILINOS:PE_TPSL_64
↪ :PE_TPSL:PE_PETSC:PE_PARALLEL_NETCDF:PE_NETCDF_H
↪ DF5PARALLEL:PE_NETCDF:PE_MPICH:PE_LIBSCI:PE_HDF5
↪ _PARALLEL:PE_HDF5:PE_GA:PE_FFTW2:PE_FFTW
MAIL=/var/mail/USER
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1 5.3
↪ 4.9
PATH=/global/homes/w/USER/.cargo/bin:/global/homes/w
↪ /USER/tools/cmake-3.14.0-rc3/bin:/global/homes/w
↪ /USER/software/cgdb/release/bin:/global/homes/w/
↪ USER/.cargo/bin:/global/homes/w/USER/tools/cmake
↪ -3.14.0-rc3/bin:/global/homes/w/USER/software/cgd
↪ b/release/bin:/usr/common/software/darshan/3.1.4
↪ /bin:/usr/common/software/altd/2.0/bin:/usr/comm
↪ on/software/bin:/usr/common/mss/bin:/usr/common/
↪ nsg/bin:/opt/java/jdk1.8.0_51/bin:/opt/cray/pe/h
↪ df5-parallel/1.10.2.0/bin:/opt/cray/pe/hdf5/1.10
↪ .2.0/bin:/opt/cray/rca/2.2.18-6.0.7.0_33.3__g2aa
↪ 4f39.ari/bin:/opt/cray/alps/6.6.43-6.0.7.0_26.4_
↪ _ga796da3.ari/sbin:/opt/cray/job/2.2.3-6.0.7.0_4
↪ 4.1__g6c4e934.ari/bin:/opt/cray/pe/craype/2.5.15
↪ /bin:/opt/gcc/7.3.0/bin:/global/homes/w/USER/.ca
↪ rgo/bin:/global/homes/w/USER/tools/cmake-3.14.0-
↪ rc3/bin:/global/homes/w/USER/software/cgdb/relea
↪ se/bin:/opt/cray/pe/mpt/7.7.3/gni/bin:/opt/ovis/
↪ bin:/opt/ovis/sbin:/usr/syscom/nsg/sbin:/usr/sys
↪ com/nsg/bin:/opt/cray/pe/modules/3.2.10.6/bin:/u
↪ sr/local/bin:/usr/bin:/bin:/usr/bin/X11:/usr/gam
↪ es:/usr/lib/mit/bin:/usr/lib/mit/sbin:/opt/cray/
↪ pe/bin
PE_MPICH_GENCOMPILERS_GNU=7.1 5.1 4.9
COLOR_GRAY=\e[0;30m
LIBGL_ALWAYS_INDIRECT=1
MODULE_VERSION=3.2.10.6

```

MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

```

PE_TPSL_DEFAULT_GENCOMPS_CRAY_interlagos=86
PE_HDF5_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /hdf5/1.10.2.0/@PRGENV@/@PE_HDF5_DEFAULT_GENCOMP
↪ S@/lib/pkgconfig
SLURM_STEP_LAUNCHER_PORT=62382
PE_NETCDF_DEFAULT_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_PARALLEL_NETCDF_DEFAULT_GENCOMPILERS_GNU=5.1 4.9
PE_FFTW_DEFAULT_TARGET_abudhabi=abudhabi
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1 5.3
↪ 4.9
PE_PETSC_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_skylake=8.2
↪ 7.1 6.1
XTPE_NETWORK_TARGET=aries
PE_HDF5_PARALLEL_BASEDIR=/opt/cray/pe/hdf5-parallel/
↪ 1.10.2.0
ATP_IGNORE_SIGTERM=1
CPU=x86_64
SLURM_WORKING_CLUSTER=edison:ct11:6817:8448
PMI_CRAY_NO_SMP_ORDER=0
_=/usr/bin/env
SLURM_JOB_ID=12930506
PE_NETCDF_HDF5PARALLEL_DEFAULT_GENCOMPS_GNU=
LDMSD_PLUGIN_LIBPATH=/opt/ovis/lib64/ovis-ldms
PE_SMA_DEFAULT_DIR_CRAY_DEFAULT64=64
PE_NETCDF_HDF5PARALLEL_DEFAULT_FIXED_PRGENV=CRAY
↪ INTEL
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
CRAY_PE_MODULES=cray-hdf5-parallel
CRAY_HDF5_DIR=/opt/cray/pe/hdf5-parallel/1.10.2.0
PE_HDF5_PARALLEL_DEFAULT_GENCOMPS_GNU=
PE_HDF5_PARALLEL_DEFAULT_FIXED_PRGENV=CRAY INTEL
JAVA_BINDIR=/opt/java/jdk1.8.0_51/bin
PMI_NO_FORK=1
CRAY_UDREG_POST_LINK_OPTS=-L/opt/cray/udreg/2.3.2-6.
↪ 0.7.0_33.18__g5196236.ari/lib64
PE_TPSL_64_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU
↪ GNU64 INTEL INTEL64
PE_TPSL_DEFAULT_GENCOMPS_INTEL_interlagos=160
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_sandybridge=86
PE_MPICH_VOLATILE_PRGENV=CRAY GNU
PWD=/global/homes/w/USER/software/MIQS/build
CRAYPE_VERSION=2.5.15
PE_TPSL_DEFAULT_GENCOMPS_GNU_mic_knl=71 53
CRAY_ALPS_POST_LINK_OPTS=-L/opt/cray/alps/6.6.43-6.0
↪ .7.0_26.4__ga796da3.ari/lib64
INPUTRC=/etc/inputrc
SLURM_STEPID=0
SLURM_JOB_USER=USER
COLOR_LIGHT_CYAN=ve[1;36m
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_sandybridge=7.1
↪ 5.3 4.9
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_MPICH_DEFAULT_GENCOMPS_CRAY=86

```

```

TARGETMODULES=craype-abudhabi:craype-abudhabi-cu:cray
↪ ype-accel-host:craype-accel-nvidia20:craype-acce
↪ l-nvidia30:craype-accel-nvidia35:craype-barcelon
↪ a:craype-broadwell:craype-haswell:craype-hugepag
↪ es128K:craype-hugepages128M:craype-hugepages16M:
↪ craype-hugepages256M:craype-hugepages2M:craype-h
↪ ugepages32M:craype-hugepages4M:craype-hugepages5
↪ 12K:craype-hugepages512M:craype-hugepages64M:cray
↪ ype-hugepages8M:craype-intel-knc:craype-interlag
↪ os:craype-interlagos-cu:craype-istanbul:craype-i
↪ vybridge:craype-mc12:craype-mc8:craype-mic-knl:c
↪ raype-network-aries:craype-network-gemini:craype
↪ -network-infiniband:craype-network-none:craype-ne
↪ twork-seastar:craype-sandybridge:craype-shanghai
↪ :craype-target-compute_node:craype-target-local_
↪ host:craype-target-native:craype-xeon:xtpe-barce
↪ lona:xtpe-interlagos:xtpe-interlagos-cu:xtpe-ist
↪ anbul:xtpe-mc12:xtpe-mc8:xtpe-network-gemini:xtp
↪ e-network-seastar:xtpe-shanghai:xtpe-target-nati
↪ ve:xtpe-xeon
HDF5_ROOT=/opt/cray/pe/hdf5-parallel/1.10.2.0/GNU/7.1
_LMFILES_=/opt/cray/pe/modulefiles/modules/3.2.10.6:
↪ /usr/syscom/nsg/modulefiles/nsg/1.2.0:/opt/cray/
↪ pe/modulefiles/cray-mpich/7.7.3:/opt/modulefiles
↪ /gcc/7.3.0:/opt/cray/pe/craype/2.5.15/modulefile
↪ s/craype-ivybridge:/opt/cray/pe/craype/2.5.15/mo
↪ dulefiles/craype-network-aries:/opt/cray/pe/modu
↪ lefiles/craype/2.5.15:/opt/cray/pe/modulefiles/c
↪ ray-libsci/18.07.1:/opt/cray/ari/modulefiles/udr
↪ eg/2.3.2-6.0.7.0_33.18__g5196236.ari:/opt/cray/a
↪ ri/modulefiles/ugni/6.0.14.0-6.0.7.0_23.1__gea11
↪ d3d.ari:/opt/cray/pe/modulefiles/pmi/5.0.14:/opt
↪ /cray/ari/modulefiles/dmapp/7.1.1-6.0.7.0_34.3__
↪ g5a674e0.ari:/opt/cray/ari/modulefiles/gni-heade
↪ rs/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari:/opt/cray
↪ /ari/modulefiles/xpmem/2.2.15-6.0.7.1_5.10__g754
↪ 9d06.ari:/opt/cray/ari/modulefiles/job/2.2.3-6.0
↪ .7.0_44.1__g6c4e934.ari:/opt/cray/ari/modulefile
↪ s/dvs/2.7.2.2.116-6.0.7.1_8.11__g0dac9eb:/opt/cr
↪ ay/ari/modulefiles/alps/6.6.43-6.0.7.0_26.4__ga7
↪ 96da3.ari:/opt/cray/ari/modulefiles/rca/2.2.18-6
↪ .0.7.0_33.3__g2aa4f39.ari:/opt/cray/pe/modulefil
↪ es/atp/2.1.3:/opt/cray/pe/modulefiles/PrgeNV-gnu
↪ /6.0.4:/opt/cray/pe/modulefiles/cray-hdf5-parall
↪ el/1.10.2.0:/opt/modulefiles/java/jdk1.8.0_51:/u
↪ sr/common/software/modulefiles/altd/2.0:/usr/com
↪ mon/software/modulefiles/darshan/3.1.4:/opt/modu
↪ lefiles/Base-opts/2.4.135-6.0.7.0_38.1__g718f891
↪ .ari
PE_LIBSCI_DEFAULT_OMP_REQUIRES=
JAVA_HOME=/opt/java/jdk1.8.0_51
SLURM_SRUN_COMM_HOST=10.128.6.142
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_skylake=8.6
PE_LIBSCI_MODULE_NAME=cray-libsci/18.07.1

```

```

SLURM_CPU_BIND_TYPE=mask_cpu:
PE_INTEL_FIXED_PKGCONFIG_PATH=/opt/cray/pe/hdf5-para
↪ ll1/1.10.2.0/INTEL/16.0/lib/pkgconfig:/opt/cray
↪ /pe/mpt/7.7.3/gni/mpich-intel/16.0/lib/pkgconfig
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_x86_skylake=82 71 61
PE_MPICH_NV_LIBS_nvidia20=-lcudart
JAVA_VERSION=jdk1.8.0_51
PYTHONSTARTUP=/etc/pythonstart
PE_MPICH_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/mpt/7.
↪ 7.3/gni/mpich-@PRGENV@PE_MPICH_DIR_DEFAULT64@/
↪ PE_MPICH_GENCOMPS@lib/pkgconfig
MODULEPATH=/opt/cray/pe/craype/2.5.15/modulefiles:/o
↪ pt/cray/pe/modulefiles:/opt/cray/modulefiles:/op
↪ t/modulefiles:/usr/common/software/modulefiles:/
↪ usr/syscom/nsg/modulefiles:/usr/common/das/modul
↪ efiles:/usr/common/ftg/modulefiles:/usr/common/g
↪ raphics/modulefiles:/usr/common/jgi/modulefiles:/
↪ /usr/common/tig/modulefiles:/opt/cray/craype/def
↪ ault/modulefiles:/opt/cray/ari/modulefiles
PE_LIBSCI_GENCOMPILERS_CRAY_x86_64=8.6
SLURM_UMASK=0022
SLURM_PTY_WIN_COL=71
TZ=US/Pacific
SDK_HOME=/usr/lib64/jvm/java
CRAYPE_USE_BUILDTOOLS=0
NSG_HOME=/usr/syscom/nsg
SHMEM_ABORT_ON_ERROR=1
MAN_POSIXLY_CORRECT=1
LOADEDMODULES=modules/3.2.10.6:nsg/1.2.0:cray-mpich/
↪ 7.7.3:gcc/7.3.0:craype-ivybridge:craype-network-
↪ aries:craype/2.5.15:cray-libsci/18.07.1:udreg/2.
↪ 3.2-6.0.7.0_33.18__g5196236.ari:ugni/6.0.14.0-6.
↪ 0.7.0_23.1__gea11d3d.ari:pmi/5.0.14:dmapp/7.1.1-
↪ 6.0.7.0_34.3__g5a674e0.ari:gni-headers/5.0.12.0-
↪ 6.0.7.0_24.1__g3b1768f.ari:xpmem/2.2.15-6.0.7.1-
↪ 5.10__g7549d06.ari:job/2.2.3-6.0.7.0_44.1__g6c4e
↪ 934.ari:dvs/2.7.2.2.116-6.0.7.1_8.11__g0dac9eb:a
↪ lps/6.6.43-6.0.7.0_26.4__ga796da3.ari:rca/2.2.18
↪ -6.0.7.0_33.3__g2aa4f39.ari:atp/2.1.3:PrgEnv-gnu/
↪ 6.0.4:cray-hdf5-parallel/1.10.2.0:java/jdk1.8.0-
↪ 51:altd/2.0:darshan/3.1.4:Base-opts/2.4.135-6.0.
↪ 7.0_38.1__g718f891.ari
SLURM_JOB_UID=74291
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_mic_knl=16.0
PE_PKG_CONFIG_PATH=/opt/cray/pe/cti/1.0.7/lib/pkgcon
↪ fig:/opt/cray/pe/cti/1.0.6/lib/pkgconfig
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_interlagos=82 71 53
↪ 49
PE_FFTW_DEFAULT_TARGET_ivybridge=ivybridge
CRAY_DMAPP_POST_LINK_OPTS=-L/opt/cray/dmapp/7.1.1-6.
↪ 0.7.0_34.3__g5a674e0.ari/lib64
PE_FFTW_DEFAULT_TARGET_x86_skylake=x86_skylake
PE_FFTW_DEFAULT_TARGET_share=share
SLURM_NODEID=0
SLURM_SUBMIT_DIR=/global/u2/w/USER/software/MIQS/bui
↪ ld
PE_LIBSCI_OMP_REQUIRES_openmp=_mp
CRAY_RCA_INCLUDE_OPTS=-I/opt/cray/rca/2.2.18-6.0.7.0
↪ _33.3__g2aa4f39.ari/include
↪ -I/opt/cray/krca/2.2.4-6.0.7.1_5.40__g8505b97.ar
↪ i/include
↪ -I/opt/cray-hss-devel/8.0.0/include
PE_HDF5_MODULE_NAME=cray-hdf5-parallel
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_x86_skylake=8.6
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_skylake=6.1
SLURM_STEP_RESV_PORTS=63186
PE_MPICH_CXX_PKGCONFIG_LIBS=mpichcxx
CRAY_MPICH_DIR=/opt/cray/pe/mpt/7.7.3/gni/mpich-gnu/
↪ 7.1
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_mic_knl=86
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
SLURM_TASK_PID=17798
PE_MPICH_PKGCONFIG_VARIABLES=PE_MPICH_NV_LIBS_@accel
↪ erator@:PE_MPICH_ALTERNATE_LIBS_@multithreaded@:
↪ PE_MPICH_ALTERNATE_LIBS_@dpm@
PE_LIBSCI_DEFAULT_GENCOMPS_INTEL_x86_64=160
SLURM_CPUS_ON_NODE=48
PE_TPSL_64_DEFAULT_GENCOMPILERS_CRAY_mic_knl=8.6
CRAY_MPICH2_VER=7.7.3
CRAY_PMI_POST_LINK_OPTS=-L/opt/cray/pe/pmi/5.0.14/li
↪ b64
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16.0
PE_MPICH_PKGCONFIG_LIBS=mpich
PE_HDF5_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_TPSL_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe
↪ /tps1/18.06.1/@PRGENV@/PE_TPSL_DEFAULT_GENCOMPS
↪ @/@PE_TPSL_DEFAULT_TARGET@/lib/pkgconfig
SLURM_PROCID=0
SLURM_JOB_NODELIST=nid0[1572-1573]
PE_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
GPG_TTY=/dev/pts/0
COLOR_BROWN=\e[0;33m
PE_PARALLEL_NETCDF_DEFAULT_FIXED_PRGENV=CRAY_INTEL
PE_MPICH_ALTERNATE_LIBS_multithreaded=_mt
PE_GA_DEFAULT_GENCOMPILERS_GNU=5.3 4.9
PE_LIBSCI_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/libsc
↪ i/18.07.1/@PRGENV@/PE_LIBSCI_GENCOMPS@/PE_LIBS
↪ CI_TARGET@/lib/pkgconfig
PE_NETCDF_HDF5PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
QT_SYSTEM_DIR=/usr/share/desktop-data
PE_HDF5_PARALLEL_DEFAULT_VOLATILE_PRGENV=GNU
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_haswell=82 71 53 49
PE_MPICH_TARGET_VAR_nvidia35=-lcudart
SHLVL=4
HOME=/global/homes/w/USER
PE_PKGCONFIG_PRODUCTS_DEFAULT=PE_PAPI
JDK_HOME=/usr/lib64/jvm/java
CRAY_LIBSCI_VERSION=18.07.1
SLURM_PTY_PORT=62381
OSTYPE=linux
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
LESS_ADVANCED_PREPROCESSOR=no

```

MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

```

SLURM_LOCALID=0
PE_PETSC_DEFAULT_GENCOMPILERS_CRAY_interlagos=8.6
PE_MPICH_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pj
↳ e/mpt/7.7.3/gni/mpich-@PRGENV@PE_MPICH_DEFAULT_
↳ DIR_DEFAULT64@/PE_MPICH_DEFAULT_GENCOMPS@/lib/pj
↳ kgconfig
ALTD_PATH=/usr/common/software/altd/2.0
PE_TPSL_DEFAULT_VOLATILE_PRGENV=CRAY CRAY64 GNU GNU64
↳ INTEL INTEL64
GREP_OPTIONS=--color=auto
SLURM_CLUSTER_NAME=edison
SLURM_JOB_CPUS_PER_NODE=48(x2)
PE_TPSL_DEFAULT_GENCOMPS_INTEL_sandybridge=160
CRAY_PMI_INCLUDE_OPTS=-I/opt/cray/pe/pmi/5.0.14/inclj
↳ ude
XCURSOR_THEME=DMZ
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_interlagos=86
LS_OPTIONS=-N --color=none -T 0
SLURM_JOB_GID=74291
SLURM_SUBMIT_HOST=mom2
WINDOWMANAGER=
PE_MPICH_DIR_CRAY_DEFAULT64=64
PRGENVMODULES=PrgEnv-cray:PrgEnv-gnu:PrgEnv-intel:Prj
↳ gEnv-pathscales:PrgEnv-pgi
ATP_MRNET_COMM_PATH=/opt/cray/pe/atp/2.1.3/libexec/a
↳ tp_mrnet_commnnode_wrapper
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_haswell=8.6
PKG_CONFIG_PATH_DEFAULT=/opt/cray/pe/papi/5.6.0.3/li
↳ b64/pkgconfig
GCC_PATH=/opt/gcc/7.3.0
CRAYPE_NETWORK_TARGET=aries
SLURM_GTIDS=0
SLURM_JOB_PARTITION=regular
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_mic_knl=7.1 5.3
PE_TPSL_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16.0
BASH_ENV=/global/homes/w/USER/.bashrc
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_interlagos=8.2 7.1
↳ 5.3 4.9
COLOR_LIGHT_GRAY=\e[0;37m
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_haswell=7.1 5.3 4.9
CRAY_LIBSCI_PREFIX_DIR=/opt/cray/pe/libsci/18.07.1/G
↳ NU/7.1/x86_64
PE_NETCDF_DEFAULT_REQUIRED_PRODUCTS=PE_HDF5
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_sandybridge=16j
↳ .0
PYTHONPATH=/opt/ovis/lib/python2.7/site-packages
PE_MPICH_NV_LIBS=
CRAY_GNI_HEADERS_INCLUDE_OPTS=-I/opt/cray/gni-headerj
↳ s/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari/include
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_haswell=8.2 7.1
↳ 5.3 4.9
MACHTYPE=x86_64-suse-linux
ALTD_SELECT_USERS=
LOGNAME=USER
COLOR_LIGHT_BLUE=\e[1;34m
PE_TPSL_DEFAULT_GENCOMPS_GNU_x86_64=82 71 53 49

```

```

LESS=-M -I -R
PE_TRILINOS_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH:PE_HDj
↳ F5_PARALLEL:PE_NETCDF_HDF5PARALLEL:PE_LIBSCI:PE_
↳ TPSL
PE_HDF5_PARALLEL_GENCOMPILERS_GNU=8.2 7.1 6.1 5.3 4.9
PE_HDF5_DEFAULT_GENCOMPS_GNU=
PE_MPICH_GENCOMPS_CRAY=86
DMAPP_ABORT_ON_ERROR=1
PE_MPICH_DEFAULT_GENCOMPILERS_CRAY=8.6
CVS_RSH=ssh
PE_LIBSCI_OMP_REQUIRES=
PE_TRILINOS_DEFAULT_GENCOMPS_GNU_x86_64=82 73 51 49
SLURM_STEP_NUM_TASKS=1
SLURM_JOB_ACCOUNT=m2621
PE_LIBSCI_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
TOOLMODULES=apprentice:apprentice2:atp:chapel:cray-lj
↳ gdb:cray-snplauncher:craypat:craypkg-gen:ddt:gdbj
↳ :iobuf:papi:perftools:perftools-lite:stat:totalvj
↳ iew:xt-craypat:xt-lgdb:xt-papi:xt-totalview
SSH_CONNECTION=172.73.67.84 53311 128.55.209.26 22
PE_TPSL_DEFAULT_GENCOMPILERS_CRAY_sandybridge=8.6
XDG_DATA_DIRS=/usr/share
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_interlagos=16.0
DVS_INCLUDE_OPTS=-I/opt/cray/dvs/2.7.2.2.116-6.0.7.1j
↳ _8.11__g0dac9eb/include
PE_MPICH_DEFAULT_GENCOMPS_GNU=71 51 49
PE_MPICH_DEFAULT_FIXED_PRGENV=INTEL
SLURM_JOB_NUM_NODES=2
PE_GA_DEFAULT_FIXED_PRGENV=CRAY INTEL
PE_LIBSCI_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/j
↳ pe/libsci/18.07.1/@PRGENV@/PE_LIBSCI_DEFAULT_GE
↳ NCOMPS@/PE_LIBSCI_DEFAULT_TARGET@/lib/pkgconfig
MODULESHOME=/opt/cray/pe/modules/3.2.10.6
PE_FFTW2_DEFAULT_REQUIRED_PRODUCTS=PE_MPICH
PE_TPSL_DEFAULT_GENCOMPILERS_GNU_sandybridge=8.2 7.1
↳ 5.3 4.9
PKG_CONFIG_PATH=/usr/common/software/darshan/3.1.4/1j
↳ ib/pkgconfig:/opt/cray/rca/2.2.18-6.0.7.0_33.3__j
↳ g2aa4f39.ari/lib64/pkgconfig:/opt/cray/alps/6.6.j
↳ 43-6.0.7.0_26.4__ga796da3.ari/lib64/pkgconfig:/o
↳ pt/cray/xpmem/2.2.15-6.0.7.1_5.10__g7549d06.ari/j
↳ lib64/pkgconfig:/opt/cray/gni-headers/5.0.12.0-6j
↳ .0.7.0_24.1__g3b1768f.ari/lib64/pkgconfig:/opt/c
↳ ray/dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari/lib64j
↳ /pkgconfig:/opt/cray/pe/pmi/5.0.14/lib64/pkgconfj
↳ ig:/opt/cray/ugni/6.0.14.0-6.0.7.0_23.1__gea11d3j
↳ d.ari/lib64/pkgconfig:/opt/cray/udreg/2.3.2-6.0.j
↳ 7.0_33.18__g5196236.ari/lib64/pkgconfig:/opt/cra
↳ y/pe/craype/2.5.15/pkg-config:/opt/cray/pe/iobufj
↳ /2.0.8/lib/pkgconfig:/opt/cray/pe/fftw/2.1.5.9/1j
↳ ib/pkgconfig:/opt/cray/pe/atp/2.1.3/lib/pkgconfig
COLOR_WHITE=\e[1;37m
PE_MPICH_NV_LIBS_nvidia35=-lcardart
LESSOPEN=lessopen.sh %s
PELOCAL_PRGENV=true

```

```

PE_PETSC_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pj
↪ e/petsc/3.8.4.0/complex/@PRGENV@/@PE_PETSC_DEFAU
↪ LT_GENCOMPS@/@PE_PETSC_DEFAULT_TARGET@/lib/pkgco
↪ nfig
COLOR_LIGHT_GREEN=\e[1;32m
LIBSCI_BASE_DIR=/opt/cray/pe/libsci/18.07.1
PE_TPSL_64_DEFAULT_GENCOMPS_INTEL_x86_64=160
SLURM_STEP_TASKS_PER_NODE=1
CRAY_NUM_COOKIES=2
PE_LIBSCI_DEFAULT_PKGCONFIG_VARIABLES=PE_LIBSCI_DEFA
↪ ULT_OMP_REQUIRES=@openmp@:PE_SCI_EXT_LIBPATH:PE_
↪ SCI_EXT_LIBNAME
PE_MPICH_NV_LIBS_nvidia60=-lcudart
PE_TPSL_64_DEFAULT_GENCOMPS_GNU_sandybridge=82 71 53
↪ 49
PE_TPSL_DEFAULT_GENCOMPS_INTEL_mic_knl=160
LIBSCI_VERSION=18.07.1
INFOPATH=/opt/gcc/7.3.0/snos/share/info
SLURM_STEP_NODELIST=nid01572
CRAY_COOKIES=4041998336,4042063872
XDG_RUNTIME_DIR=/run/user/74291
PE_LIBSCI_GENCOMPILERS_INTEL_x86_64=16.0
CRAY_PRE_COMPILE_OPTS=-hnetwork=aries
CRAY_ALPS_INCLUDE_OPTS=-I/opt/cray/alps/6.6.43-6.0.7
↪ .0_26.4__ga796da3.ari/include
CRAY_CPU_TARGET=sandybridge
PE_FFTW_DEFAULT_TARGET_broadwell=broadwell
NERSC_HOST=edison
PE_TPSL_64_DEFAULT_GENCOMPILERS_GNU_x86_64=8.2 7.1
↪ 5.3 4.9
MPICH_GNI_ROUTING_MODE=ADAPTIVE_1
CRAY_UGNI_INCLUDE_OPTS=-I/opt/cray/ugni/6.0.14-6.0
↪ .7.0_23.1__gea11d3d.ari/include
PE_PAPI_DEFAULT_ACCELL_FAMILY_LIBS=
PE_TPSL_64_DEFAULT_GENCOMPS_CRAY_x86_64=86
PE_LIBSCI_REQUIRED_PRODUCTS=PE_MPICH
CRAY_XPMEM_INCLUDE_OPTS=-I/opt/cray/xpmem/2.2.15-6.0
↪ .7.1_5.10__g7549d06.ari/include
craype_already_loaded=0
PE_HDF5_PARALLEL_REQUIRED_PRODUCTS=PE_MPICH
PE_TPSL_DEFAULT_GENCOMPS_INTEL_haswell=160
PE_LIBSCI_DEFAULT_GENCOMPILERS_GNU_x86_64=7.1 6.1 5.1
↪ 4.9
PE_LIBSCI_GENCOMPS_GNU_x86_64=71 61 51 49
SLURM_CPU_BIND=quiet,mask_cpu:0xFFFFFFFFFFFF
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_x86_64=16.0
PE_FFTW_DEFAULT_TARGET_x86_64=x86_64
LESSCLOSE=lessclose.sh %s %s
ATP_HOME=/opt/cray/pe/atp/2.1.3
PE_HDF5_PARALLEL_GENCOMPS_GNU=
SLURM_MEM_PER_NODE=57344
SCRATCH=/scratch2/scratchdirs/USER
PE_TPSL_64_DEFAULT_GENCOMPILERS_INTEL_x86_skylake=16
↪ .0
ALTD_WORKDIR=/scratch1/scratchdirs/altdlogs
PE_FFTW_DEFAULT_TARGET_haswell=haswell
PE_PETSC_DEFAULT_GENCOMPILERS_INTEL_haswell=16.0
PE_GA_DEFAULT_VOLATILE_PKGCONFIG_PATH=/opt/cray/pe/g
↪ a/5.3.0.8/@PRGENV@/@PE_GA_DEFAULT_GENCOMPS@/lib/
↪ pkgconfig
G_BROKEN_FILENAMES=1
PE_GA_DEFAULT_GENCOMPS_GNU=53 49
PE_PETSC_DEFAULT_GENCOMPILERS_GNU_interlagos=7.1 5.3
↪ 4.9
PE_PAPI_DEFAULT_ACCEL_LIBS=
PE_INTEL_DEFAULT_FIXED_PKGCONFIG_PATH=/opt/cray/pe/pj
↪ arallel-netcdf/1.8.1.3/INTEL/16.0/lib/pkgconfig:
↪ /opt/cray/pe/netcdf-hdf5parallel/4.6.1.3/INTEL/1
↪ 6.0/lib/pkgconfig:/opt/cray/pe/netcdf/4.6.1.3/IN
↪ TEL/16.0/lib/pkgconfig:/opt/cray/pe/mpt/7.7.3/gn
↪ i/mpich-intel/16.0/lib/pkgconfig:/opt/cray/pe/hd
↪ f5-parallel/1.10.2.0/INTEL/16.0/lib/pkgconfig:/o
↪ pt/cray/pe/hdf5/1.10.2.0/INTEL/16.0/lib/pkgconfi
↪ g:/opt/cray/pe/ga/5.3.0.8/INTEL/18.0/lib/pkgconf
↪ ig
CRAY_LD_LIBRARY_PATH=/opt/cray/pe/hdf5-parallel/1.10
↪ .2.0/GNU/7.1/lib:/opt/cray/rca/2.2.18-6.0.7.0_33
↪ .3__g2aa4f39.ari/lib64:/opt/cray/alps/6.6.43-6.0
↪ .7.0_26.4__ga796da3.ari/lib64:/opt/cray/xpmem/2.
↪ 2.15-6.0.7.1_5.10__g7549d06.ari/lib64:/opt/cray/
↪ dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari/lib64:/op
↪ t/cray/pe/pmi/5.0.14/lib64:/opt/cray/ugni/6.0.14
↪ .0-6.0.7.0_23.1__gea11d3d.ari/lib64:/opt/cray/ud
↪ reg/2.3.2-6.0.7.0_33.18__g5196236.ari/lib64:/opt
↪ /cray/pe/libsci/18.07.1/GNU/7.1/x86_64/lib:/opt/
↪ cray/pe/mpt/7.7.3/gni/mpich-gnu/7.1/lib
PE_SMA_DEFAULT_DIR_PGI_DEFAULT64=64
COLORTERM=1
PE_PETSC_DEFAULT_GENCOMPS_CRAY_haswell=86
PE_PETSC_DEFAULT_GENCOMPS_GNU_x86_64=71 53 49
PE_PETSC_DEFAULT_GENCOMPS_INTEL_x86_64=160
PE_MPICH_DEFAULT_DIR_CRAY_DEFAULT64=64
JAVA_ROOT=/opt/java/jdk1.8.0_51
BASH_FUNC_module%=( ) { eval
↪ ` /opt/cray/pe/modules/3.2.10.6/bin/modulecmd
↪ bash $* `
}
+ lsb_release -a
LSB Version: n/a
Distributor ID: SUSE
Description: SUSE Linux Enterprise Server 12
↪ SP3
Release: 12.3
Codename: n/a
+ uname -a
Linux nid01572 4.4.103-6.38_4.0.153-cray_ari_c #1 SMP
↪ Thu Nov 1 16:05:05 UTC 2018 (6ef8fef) x86_64
↪ x86_64 x86_64 GNU/Linux
+ lscpu
Architecture: x86_64
CPU op-mode(s): 32-bit, 64-bit
Byte Order: Little Endian

```

MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

```

CPU(s): 48
On-line CPU(s) list: 0-47
Thread(s) per core: 2
Core(s) per socket: 12
Socket(s): 2
NUMA node(s): 2
Vendor ID: GenuineIntel
CPU family: 6
Model: 62
Model name: Intel(R) Xeon(R) CPU E5-2695 v2
↳ @ 2.40GHz
Stepping: 4
CPU MHz: 2401.000
CPU max MHz: 2401.0000
CPU min MHz: 1200.0000
BogoMIPS: 4799.59
Virtualization: VT-x
L1d cache: 32K
L1i cache: 32K
L2 cache: 256K
L3 cache: 30720K
NUMA node0 CPU(s): 0-11,24-35
NUMA node1 CPU(s): 12-23,36-47
Flags: fpu vme de pse tsc msr pae mce
↳ cx8 apic sep mtrr pge mca cmov pat pse36 clflush
↳ dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx
↳ pdpe1gb rdtscp lm constant_tsc arch_perfmon pbs
↳ bts rep_good nopl xtopology nonstop_tsc aperfmperf
↳ eagerfpu pni pclmulqdq dtes64 monitor ds_cpl vmx
↳ smx est tm2 ssse3 cx16 xtpr pdcm pcid dca sse4_1
↳ sse4_2 x2apic popcnt tsc_deadline_timer aes xsave
↳ avx f16c rdrand lahf_lm ida arat epb pln pts
↳ dtherm spec_ctrl kaiser tpr_shadow vnmi
↳ flexpriority ept vpid fsgsbase smep erms xsaveopt
+ cat /proc/meminfo
MemTotal: 65938820 kB
MemFree: 62827876 kB
MemAvailable: 62558036 kB
Buffers: 4988 kB
Cached: 268892 kB
SwapCached: 0 kB
Active: 163620 kB
Inactive: 229076 kB
Active(anon): 148272 kB
Inactive(anon): 205596 kB
Active(file): 15348 kB
Inactive(file): 23480 kB
Unevictable: 2576 kB
Mlocked: 2576 kB
SwapTotal: 0 kB
SwapFree: 0 kB
Dirty: 0 kB
Writeback: 0 kB
AnonPages: 121548 kB
Mapped: 57736 kB
Shmem: 234884 kB

```

```

Slab: 382700 kB
SReclaimable: 22840 kB
SUnreclaim: 359860 kB
KernelStack: 13584 kB
PageTables: 4656 kB
NFS_Unstable: 0 kB
Bounce: 0 kB
WritebackTmp: 0 kB
CommitLimit: 32969408 kB
Committed_AS: 559876 kB
VmallocTotal: 34359738367 kB
VmallocUsed: 0 kB
VmallocChunk: 0 kB
HardwareCorrupted: 0 kB
HugePages_Total: 0
HugePages_Free: 0
HugePages_Rsvd: 0
HugePages_Surp: 0
Hugepagesize: 2048 kB
DirectMap4k: 12028 kB
DirectMap2M: 2048000 kB
DirectMap1G: 67108864 kB
+ inxi -F -c0
sc_env_info.sh: line 14: inxi: command not found
+ lsblk -a
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
loop0 7:0 0 128K 0 loop
↳ /var/opt/cray/imps-distribution/squash
loop1 7:1 0 1.1M 0 loop
↳ /var/opt/cray/imps-distribution/squash
loop2 7:2 0 2.4G 1 loop /.rootfs_lower_ro
loop3 7:3 0 23.8G 1 loop
↳ /var/opt/cray/imps-image-binding/PE_x8
loop4 7:4 0 0 loop
loop5 7:5 0 0 loop
loop6 7:6 0 0 loop
loop7 7:7 0 0 loop
loop8 7:8 0 0 loop
loop9 7:9 0 0 loop
loop10 7:10 0 0 loop
loop11 7:11 0 0 loop
loop12 7:12 0 0 loop
loop13 7:13 0 0 loop
loop14 7:14 0 0 loop
loop15 7:15 0 0 loop
loop16 7:16 0 0 loop
loop17 7:17 0 0 loop
loop18 7:18 0 0 loop
loop19 7:19 0 0 loop
loop20 7:20 0 0 loop
loop21 7:21 0 0 loop
loop22 7:22 0 0 loop
loop23 7:23 0 0 loop
loop24 7:24 0 0 loop
loop25 7:25 0 0 loop
loop26 7:26 0 0 loop

```

```

loop27 7:27 0 0 loop
loop28 7:28 0 0 loop
loop29 7:29 0 0 loop
loop30 7:30 0 0 loop
loop31 7:31 0 0 loop
loop32 7:32 0 0 loop
loop33 7:33 0 0 loop
loop34 7:34 0 0 loop
loop35 7:35 0 0 loop
loop36 7:36 0 0 loop
loop37 7:37 0 0 loop
loop38 7:38 0 0 loop
loop39 7:39 0 0 loop
loop40 7:40 0 0 loop
loop41 7:41 0 0 loop
loop42 7:42 0 0 loop
loop43 7:43 0 0 loop
loop44 7:44 0 0 loop
loop45 7:45 0 0 loop
loop46 7:46 0 0 loop
loop47 7:47 0 0 loop
loop48 7:48 0 0 loop
loop49 7:49 0 0 loop
loop50 7:50 0 0 loop
loop51 7:51 0 0 loop
loop52 7:52 0 0 loop
loop53 7:53 0 0 loop
loop54 7:54 0 0 loop
loop55 7:55 0 0 loop
loop56 7:56 0 0 loop
loop57 7:57 0 0 loop
loop58 7:58 0 0 loop
loop59 7:59 0 0 loop
loop60 7:60 0 0 loop
loop61 7:61 0 0 loop
loop62 7:62 0 0 loop
loop63 7:63 0 0 loop
loop64 7:64 0 0 loop
loop65 7:65 0 0 loop
loop66 7:66 0 0 loop
loop67 7:67 0 0 loop
loop68 7:68 0 0 loop
loop69 7:69 0 0 loop
loop70 7:70 0 0 loop
loop71 7:71 0 0 loop
loop72 7:72 0 0 loop
loop73 7:73 0 0 loop
loop74 7:74 0 0 loop
loop75 7:75 0 0 loop
loop76 7:76 0 0 loop
loop77 7:77 0 0 loop
loop78 7:78 0 0 loop
loop79 7:79 0 0 loop
loop80 7:80 0 0 loop
loop81 7:81 0 0 loop
loop82 7:82 0 0 loop
loop83 7:83 0 0 loop

```

```

loop84 7:84 0 0 loop
loop85 7:85 0 0 loop
loop86 7:86 0 0 loop
loop87 7:87 0 0 loop
loop88 7:88 0 0 loop
loop89 7:89 0 0 loop
loop90 7:90 0 0 loop
loop91 7:91 0 0 loop
loop92 7:92 0 0 loop
loop93 7:93 0 0 loop
loop94 7:94 0 0 loop
loop95 7:95 0 0 loop
loop96 7:96 0 0 loop
loop97 7:97 0 0 loop
loop98 7:98 0 0 loop
loop99 7:99 0 0 loop
loop100 7:100 0 0 loop
loop101 7:101 0 0 loop
loop102 7:102 0 0 loop
loop103 7:103 0 0 loop
loop104 7:104 0 0 loop
loop105 7:105 0 0 loop
loop106 7:106 0 0 loop
loop107 7:107 0 0 loop
loop108 7:108 0 0 loop
loop109 7:109 0 0 loop
loop110 7:110 0 0 loop
loop111 7:111 0 0 loop
loop112 7:112 0 0 loop
loop113 7:113 0 0 loop
loop114 7:114 0 0 loop
loop115 7:115 0 0 loop
loop116 7:116 0 0 loop
loop117 7:117 0 0 loop
loop118 7:118 0 0 loop
loop119 7:119 0 0 loop
loop120 7:120 0 0 loop
loop121 7:121 0 0 loop
loop122 7:122 0 0 loop
loop123 7:123 0 0 loop
loop124 7:124 0 0 loop
loop125 7:125 0 0 loop
loop126 7:126 0 0 loop
loop127 7:127 0 0 loop

```

```

+ lsscsi -s
+ module list
++ /opt/cray/pe/modules/3.2.10.6/bin/modulecmd bash
↪ list

```

Currently Loaded Modulefiles:

- 1) modules/3.2.10.6
- 2) cray-mpich/7.7.3
- 3) gcc/7.3.0
- 4) craype-ivybridge
- 5) craype-network-aries
- 6) craype/2.5.15
- 7) cray-libsci/18.07.1

MIQS: Metadata Indexing and Querying Service for Self-Describing File Formats

```
8) udreg/2.3.2-6.0.7.0_33.18__g5196236.ari
9) ugni/6.0.14.0-6.0.7.0_23.1__gea11d3d.ari
10) pmi/5.0.14
11) dmapp/7.1.1-6.0.7.0_34.3__g5a674e0.ari
12) gni-headers/5.0.12.0-6.0.7.0_24.1__g3b1768f.ari
13) xpmem/2.2.15-6.0.7.1_5.10__g7549d06.ari
14) job/2.2.3-6.0.7.0_44.1__g6c4e934.ari
15) dvs/2.7_2.2.116-6.0.7.1_8.11__g0dac9eb
16) alps/6.6.43-6.0.7.0_26.4__ga796da3.ari
17) rca/2.2.18-6.0.7.0_33.3__g2aa4f39.ari
18) atp/2.1.3
19) PrgEnv-gnu/6.0.4
20) cray-hdf5-parallel/1.10.2.0
21) java/jdk1.8.0_51
22) altd/2.0
23) darshan/3.1.4
24) Base-opts/2.4.135-6.0.7.0_38.1__g718f891.ari
+ eval
+ nvidia-smi
NVIDIA-SMI has failed because it couldn't communicate
↳ with the NVIDIA driver. Make sure that the latest
↳ NVIDIA driver is installed and running.

+ lshw -short -quiet -sanitize
+ cat
sc_env_info.sh: line 19: lshw: command not found
+ lspci
sc_env_info.sh: line 19: lspci: command not found
```