

---

# **ADIOS2 Documentation**

***Release 2.10.0-rc1***

**William F Godoy**

**Apr 26, 2024**



# INTRODUCTION

<b>1</b>	<b>What's new in 2.10?</b>	<b>3</b>
1.1	Python . . . . .	3
1.2	New/updated features . . . . .	3
1.3	Packaging . . . . .	4
<b>2</b>	<b>What's new in 2.9?</b>	<b>5</b>
2.1	General . . . . .	5
2.2	File I/O . . . . .	5
2.3	Staging . . . . .	6
2.4	Experimental features . . . . .	6
<b>3</b>	<b>Introduction</b>	<b>7</b>
3.1	What ADIOS2 is and isn't . . . . .	7
3.2	Adaptable IO beyond files in Scientific Data Lifecycles . . . . .	8
<b>4</b>	<b>Install from Source</b>	<b>9</b>
4.1	Building, Testing, and Installing ADIOS 2 . . . . .	9
4.2	CMake Options . . . . .	11
4.3	Building on HPC Systems . . . . .	12
4.4	Installing the ADIOS2 library and the C++ and C bindings . . . . .	14
4.5	Enabling the Python bindings . . . . .	14
4.6	Enabling the Fortran bindings . . . . .	15
4.7	Running Tests . . . . .	15
4.8	Running Examples . . . . .	16
<b>5</b>	<b>As Package</b>	<b>17</b>
5.1	Conda . . . . .	17
5.2	PyPI . . . . .	17
5.3	Spack . . . . .	17
5.4	Docker . . . . .	17
<b>6</b>	<b>Linking ADIOS 2</b>	<b>19</b>
6.1	From CMake . . . . .	19
6.2	From non-CMake build systems . . . . .	19
<b>7</b>	<b>Use on DOE machines</b>	<b>21</b>
7.1	NERSC Perlmutter . . . . .	21
7.2	OLCF Frontier . . . . .	22
7.3	ALCF Aurora . . . . .	23
<b>8</b>	<b>Interface Components</b>	<b>25</b>

8.1	Components Overview . . . . .	25
8.2	ADIOS . . . . .	27
8.3	IO . . . . .	28
8.4	Variable . . . . .	32
8.5	Attribute . . . . .	39
8.6	Engine . . . . .	40
8.7	Operator . . . . .	49
8.8	Runtime Configuration Files . . . . .	51
8.9	Anatomy of an ADIOS Program . . . . .	53
<b>9</b>	<b>Supported Virtual Engine Names</b>	<b>57</b>
9.1	Virtual Engine Setups . . . . .	58
<b>10</b>	<b>Supported Engines</b>	<b>59</b>
10.1	BP5 . . . . .	59
10.2	BP4 . . . . .	64
10.3	BP3 . . . . .	66
10.4	HDF5 . . . . .	68
10.5	SST Sustainable Staging Transport . . . . .	69
10.6	SSC Strong Staging Coupler . . . . .	72
10.7	DataMan for Wide Area Network Data Staging . . . . .	73
10.8	DataSpaces . . . . .	74
10.9	Inline for zero-copy . . . . .	75
10.10	Null . . . . .	76
10.11	Plugin Engine . . . . .	76
<b>11</b>	<b>Supported Operators</b>	<b>77</b>
11.1	CompressorZFP . . . . .	77
11.2	Plugin Operator . . . . .	78
11.3	Encryption . . . . .	78
<b>12</b>	<b>Full APIs</b>	<b>79</b>
12.1	C++11 bindings . . . . .	79
12.2	Fortran bindings . . . . .	104
12.3	C bindings . . . . .	129
<b>13</b>	<b>High-Level APIs</b>	<b>153</b>
13.1	C++ High-Level API . . . . .	154
13.2	Matlab simple bindings . . . . .	164
<b>14</b>	<b>Python APIs</b>	<b>167</b>
14.1	Python Example Code . . . . .	167
14.2	adios2 classes for Python . . . . .	171
14.3	Python bindings to C++ . . . . .	184
14.4	Transition from old API to new API . . . . .	188
<b>15</b>	<b>Aggregation</b>	<b>191</b>
15.1	Aggregation in BP5 . . . . .	191
<b>16</b>	<b>Memory Management</b>	<b>193</b>
16.1	BP4 buffering . . . . .	193
16.2	BP5 buffering . . . . .	193
16.3	Span object in internal buffer . . . . .	194
<b>17</b>	<b>GPU-aware I/O</b>	<b>195</b>

17.1	Building ADIOS2 with a GPU backend . . . . .	195
17.2	Writing GPU buffers . . . . .	196
17.3	Reading GPU buffers . . . . .	197
17.4	Build scripts . . . . .	198
<b>18</b>	<b>ADIOS2 query API</b>	<b>199</b>
18.1	The interface . . . . .	199
18.2	Code EXAMPLES: . . . . .	200
<b>19</b>	<b>Plugins</b>	<b>201</b>
19.1	Writing Your Plugin Class . . . . .	201
19.2	Build Shared Library . . . . .	203
19.3	Using Your Plugin in an Application . . . . .	204
<b>20</b>	<b>Campaign Management</b>	<b>207</b>
20.1	The idea . . . . .	207
20.2	Setup . . . . .	209
20.3	Remote access . . . . .	210
20.4	Requirements . . . . .	211
20.5	Limitations . . . . .	211
<b>21</b>	<b>ADIOS2 in ECP hardware</b>	<b>213</b>
21.1	OLCF CRUSHER . . . . .	213
<b>22</b>	<b>Overview</b>	<b>215</b>
<b>23</b>	<b>Download And Build</b>	<b>217</b>
<b>24</b>	<b>Basic Tutorials</b>	<b>219</b>
24.1	Hello World . . . . .	219
24.2	Variables . . . . .	224
24.3	Attributes . . . . .	233
24.4	Operators . . . . .	240
24.5	Steps . . . . .	246
<b>25</b>	<b>HDF5 API Support through VOL</b>	<b>253</b>
25.1	Disclaimer . . . . .	253
25.2	External . . . . .	253
25.3	Internal . . . . .	253
<b>26</b>	<b>Command Line Utilities</b>	<b>255</b>
26.1	bpls : Inspecting Data . . . . .	255
26.2	adios_reorganize . . . . .	258
26.3	adios2-config . . . . .	261
26.4	sst_conn_tool : SST network connectivity tool . . . . .	261
<b>27</b>	<b>Visualizing Data</b>	<b>263</b>
27.1	Using VTK and Paraview . . . . .	263
<b>28</b>	<b>FAQ</b>	<b>267</b>
28.1	MPI vs Non-MPI . . . . .	267
28.2	APIs . . . . .	267
28.3	Building on Titan . . . . .	267
28.4	Building and Running on Fujitsu FX100 . . . . .	267
28.5	FAQs Answered . . . . .	268

<b>29 Advice</b>	<b>273</b>
<b>Index</b>	<b>275</b>

Funded by the [Exascale Computing Project \(ECP\)](#), U.S. Department of Energy





## WHAT'S NEW IN 2.10?

This is a major release with new features and lots of bug fixes. The main new feature is the new Python API.

### 1.1 Python

Before, ADIOS had two separate APIs for Python. The low-level (“Full”) API was written with Pybind11 and directly mimicked the C++ API. The high-level API was another, smaller, and more pythonesque API that allowed for easier scripting with Python. The main problems with these two were that they were independent, and that the high-level API was not complete. Once a developer needed a feature only available in the full API, they had to start from scratch writing a script with the full API.

In 2.10, there is officially one Python API, written in Python, which in turn uses the old Pybind11 classes. The new API combines the high-level features of the old high-level API – hopefully in a more consistent and likeable way, – and the full feature set of the low-level bindings.

---

**Note:** Old scripts that used the full API can still run without almost any modification, just change the import line from `import adios2` to `import adios2.bindings as adios2`

Old scripts that used the high-level API must be modified to make them work with the new API, see [Transition from old API to new API](#)

---

See [Python API](#)

### 1.2 New/updated features

- BP5 is supported on Windows now
- SST and DataMan staging engines are GPU-Aware now
- SYCL support added for Intel GPUs (besides CUDA and HIP for NVidia and AMD GPUs)
- the SST/libfabric data transport now works on Frontier (besides the MPI data transport)

## 1.3 Packaging

- adios2 package is now on [PyPi](#)

## WHAT'S NEW IN 2.9?

This is a major release with new features and lots of bug fixes.

### 2.1 General

- GPU-Aware I/O enabled by using Kokkos. Device pointers can be passed to Put()/Get() calls directly. Kokkos 3.7.x required for this release. Works with CUDA, HIP and Kokkos applications. [https://adios2.readthedocs.io/en/latest/advanced/gpu\\_aware.html#gpu-aware-i-o](https://adios2.readthedocs.io/en/latest/advanced/gpu_aware.html#gpu-aware-i-o)
- GPU-compression. MGARD and ZFP operators can compress data on GPU if they are built for GPU. MGARD operator can be fed with host/device pointers and will move data automatically. ZFP operator requires matching data and compressor location.
- Joined Array concept (besides Global Array and Local Array), which lets writers dump Local Arrays (no offsets no global shape) that are put together into a Global Array by the reader. One dimension of the arrays is selected for this join operation, while other dimensions must be the same for all writers. <https://adios2.readthedocs.io/en/latest/components/components.html?highlight=Joined#shapes>

### 2.2 File I/O

- Default File engine is now BP5. If for some reason this causes problems, manually specify using “BP4” for your application.
- BP5 engine supports multithreaded reading to accelerate read performance for low-core counts.
- BP5 Two level metadata aggregation and reduction reduced memory impact of collecting metadata and therefore is more scalable in terms of numbers of variables and writers than BP4.
- Uses Blosc-2 instead of Blosc for lossless compression. The new compression operator is backward compatible with old files compressed with blosc. The name of the operator remains “blosc”.

## 2.3 Staging

- UCX dataplane added for SST staging engine to support networks under the UCX consortium
- MPI dataplane added for SST staging engine. It relies on MPI intercommunicators to connect multiple independent MPI applications for staging purposes. Applications must enable multithreaded MPI for this dataplane.

## 2.4 Experimental features

- Preliminary support for data structs. A struct can have single variables of basic types, and 1D fixed size arrays of basic types. Supported by BP5, SST and SSC engines.

## INTRODUCTION

**ADIOS2** is the latest implementation of the [Adaptable Input Output System](#). This brand new architecture continues the performance legacy of ADIOS1, and extends its capabilities to address the extreme challenges of scientific data IO.

The [ADIOS2 repo](#) is hosted at [GitHub](#).

The ADIOS2 infrastructure is developed as a multi-institutional collaboration between

- [Oak Ridge National Laboratory](#)
- [Kitware Inc.](#)
- [Lawrence Berkeley National Laboratory](#)
- [Georgia Institute of Technology](#)
- [Rutgers University](#)

The key aspects ADIOS2 are

1. **Modular architecture:** ADIOS2 takes advantage of the major features of C++11. The architecture utilizes a balanced combination of runtime polymorphism and template meta-programming to expose intuitive abstractions for a broad range of IO applications.
2. **Community:** By maintaining coding standards, collaborative workflows, and understandable documentation, ADIOS2 lowers the barriers to entry for scientists to meaningfully interact with the code.
3. **Sustainability:** Continuous integration and unit testing ensure that ADIOS2 evolves responsibly. Bug reports are triaged and fixed in a timely manner and can be reported on [GitHub](#).
4. **Language Support:** In addition to the native C++, bindings for Python, C, Fortran and Matlab are provided.
5. **Commitment:** ADIOS2 is committed to the HPC community, releasing a new version every 6 months.

*ADIOS2 is funded by the Department of Energy as part of the [Exascale Computing Project](#).*

### 3.1 What ADIOS2 is and isn't

**ADIOS2 is:**

- **A Unified High-performance I/O Framework:** using the same abstraction API ADIOS2 can transport and transform groups of self-describing data variables and attributes across different media (file, wide-area-network, in-memory staging, etc.) with performance an ease of use as the main goals.
- **MPI-based:** parallel MPI applications as well as serial codes can use it
- **Streaming-oriented:** ADIOS2 favors codes transferring a group of variables asynchronously wherever possible. Moving one variable at a time, in synchronous fashion, is the special case rather than normal.

- **Step-based:** to resemble actual production of data in “steps” of variable groups, for either streaming or random-access (file) media
- **Free and open-source:** ADIOS2 is permissibly licensed under the OSI-approved Apache 2 license.
- **Extreme scale I/O:** ADIOS2 is being used in supercomputer applications that write and read up to several petabytes in a single simulation run. ADIOS2 is designed to provide scalable I/O on the largest supercomputers in the world.

**ADIOS2 is not:**

- **A file-only I/O library:** Code coupling and in situ analysis is possible through files but special engines are available to achieve the same thing faster through TCP, RDMA and MPI communication. High performance write/read using a file system is a primary goal of ADIOS2 though.
- **MPI-only**
- **A Hierarchical Model:** Data hierarchies can be built on top of the ADIOS2 according to the application, but ADIOS2 sits a layer of abstraction beneath this.
- **A Memory Manager Library:** we don’t own or manage the application’s memory

## 3.2 Adaptable IO beyond files in Scientific Data Lifecycles

Performant and usable tools for data management at scale are essential in an era where scientific breakthroughs are collaborative, multidisciplinary, and computational. ADIOS2 is an *adaptable*, *scalable*, and *unified* framework to aid scientific applications when data transfer volumes exceed the capabilities of traditional file I/O.

ADIOS2 provides

- Custom application management of massive data sets, starting from generation, analysis, and movement, as well as short-term and long-term storage.
- Self-describing data in binary-packed (*.bp*) format for rapid metadata extraction
- An ability to separate and extract relevant information from large data sets
- The capability to make real-time decisions based on in-transit or in-situ analytics
- The ability to expand to other transport mechanisms such wide area networks, remote direct memory access, and shared memory, with minimal overhead
- The ability to utilize the full capabilities of emergent hardware technologies, such as high-bandwidth memory and burst buffers

## INSTALL FROM SOURCE

ADIOS2 uses **CMake** for building, testing and installing the library and utilities.

### 4.1 Building, Testing, and Installing ADIOS 2

To build ADIOS v2.x, clone the repository and invoke the canonical CMake build sequence:

```
$ git clone https://github.com/ornladios/ADIOS2.git ADIOS2
$ mkdir adios2-build && cd adios2-build
$ cmake ../ADIOS2 cmake -DADIOS2_BUILD_EXAMPLES=ON
-- The C compiler identification is GNU 9.4.0
-- The CXX compiler identification is GNU 9.4.0
...
```

ADIOS2 build configuration:

ADIOS Version: 2.10.0

C++ Compiler : GNU 9.4.0  
/usr/bin/c++

Fortran Compiler : GNU 9.4.0  
/usr/bin/f95

Installation prefix: /usr/local  
bin: bin  
lib: lib  
include: include  
cmake: lib/cmake/adios2  
python: lib/python3/dist-packages

...

Features:

Library Type: shared  
Build Type: Release  
Testing: OFF  
Examples: ON

Build Options:

DataMan	: ON
DataSpaces	: OFF
HDF5	: OFF
HDF5_VOL	: OFF

(continues on next page)

(continued from previous page)

```

MHS                : ON
SST                : ON
Fortran            : ON
MPI                : ON
Python             : ON
PIP                : OFF
Blosc2             : OFF
BZip2              : ON
LIBPRESSIO         : OFF
MGARD              : OFF
MGARD_MDR          : OFF
PNG                : OFF
SZ                 : OFF
ZFP                : ON
DAOS               : OFF
IME                : OFF
O_DIRECT           : ON
Sodium             : ON
Catalyst           : OFF
SysVShMem          : ON
UCX                : OFF
ZeroMQ             : ON
Profiling          : ON
Endian_Reverse     : OFF
Derived_Variable   : OFF
AWSSDK             : OFF
GPU_Support        : OFF
CUDA               : OFF
Kokkos             : OFF
Kokkos_CUDA        : OFF
Kokkos_HIP         : OFF
Kokkos_SYCL        : OFF
Campaign           : OFF

```

If a desired feature is OFF in the report above, tell cmake where to find the required dependencies for that feature and manually turn it on. E.g.:

```
$ cmake ... -DADIOS2_USE_Blosc2=ON -DCMAKE_PREFIX_PATH="<path to c-blosc2 installation>"
```

Then compile using

```
$ make -j 16
```

Optionally, run the tests (need to configure with `-DBUILD_TESTING=ON` cmake flag)

```

$ ctest
Test project /home/wgodoy/workspace/build
   Start    1: ADIOSInterfaceWriteTest.DefineVar_int8_t_1x10
 1/295 Test #1: ADIOSInterfaceWriteTest.DefineVar_int8_t_1x10 .....
↪. Passed  0.16 sec
   Start    2: ADIOSInterfaceWriteTest.DefineVar_int16_t_1x10
 2/295 Test #2: ADIOSInterfaceWriteTest.DefineVar_int16_t_1x10 .....

```

(continues on next page)



(continued from previous page)

```

→. Passed 0.06 sec
    Start 3: ADIOSInterfaceWriteTest.DefineVar_int32_t_1x10

...

    Start 294: ADIOSBZip2Wrapper.WrongParameterValue
294/295 Test #294: ADIOSBZip2Wrapper.WrongParameterValue .....
→. Passed 0.00 sec
    Start 295: ADIOSBZip2Wrapper.WrongBZip2Name
295/295 Test #295: ADIOSBZip2Wrapper.WrongBZip2Name .....
→. Passed 0.00 sec

100% tests passed, 0 tests failed out of 295

Total Test time (real) = 95.95 sec

```

And finally, use the standard invocation to install:

```
$ make install
```

## 4.2 CMake Options

The following options can be specified with CMake's `-DVAR=VALUE` syntax. The default option is highlighted.

VAR	VAL	Description
ADIOS2_USE_	ON/C	MPI or non-MPI (serial) build.
ADIOS2_USE_	ON/C	<b>ZeroMQ</b> for the DataMan engine.
ADIOS2_USE_	ON/C	<b>HDF5</b> engine. If HDF5 is not on the syspath, it can be set using <code>-DHDF5_ROOT=/path/to/hdf5</code>
ADIOS2_USE_	ON/C	Python bindings. Python 3 will be used if found. If you want to specify a particular python version use <code>-DPYTHON_EXECUTABLE=/path/to/interpreter/python</code> <code>-DPython_FIND_STRATEGY=LOCATION</code>
ADIOS2_USE_	ON/C	Bindings for Fortran 90 or above.
ADIOS2_USE_	ON/C	Simplified Staging Engine (SST) and its dependencies, requires MPI. Can optionally use LibFabric/UCX for RDMA transport. You can specify the LibFabric/UCX path manually with the <code>-DLIBFABRIC_ROOT=...</code> or <code>-DUCX_ROOT=...</code> option.
ADIOS2_USE_	ON/C	<b>BZIP2</b> compression.
ADIOS2_USE_	ON/C	<b>ZFP</b> compression (experimental).
ADIOS2_USE_	ON/C	<b>SZ</b> compression (experimental).
ADIOS2_USE_	ON/C	<b>MGARD</b> compression (experimental).
ADIOS2_USE_	ON/C	<b>PNG</b> compression (experimental).
ADIOS2_USE_	ON/C	<b>Blosc</b> compression (experimental).
ADIOS2_USE_	ON/C	Enable endian conversion if a different endianness is detected between write and read.
ADIOS2_USE_	ON/C	DDN IME transport.

In addition to the `ADIOS2_USE_Feature` options, the following options are also available to control how the library gets built:

CMake VAR Options	Values	Description
BUILD_SHARED_LIBS	ON/OFF	Build shared libraries.
ADIOS2_BUILD_EXAMPLES	ON/OFF	Build examples.
BUILD_TESTING	ON/OFF	Build test code.
CMAKE_INSTALL_PREFIX	/path/to/install (/usr/local)	Installation location.
CMAKE_BUILD_TYPE	Debug/ <b>Release</b> /RelWithDebInfo/MinSizeRel	Compiler optimization levels.
CMAKE_PREFIX_PATH	Semi-colon separated list of paths	Location of extra dependencies

Example: Enable Fortran, disable Python bindings and ZeroMQ functionality

```
$ cmake -DADIOS2_USE_Fortran=ON -DADIOS2_USE_Python=OFF -DADIOS2_USE_ZeroMQ=OFF ../ADIOS2
```

Notes:

To provide search paths to CMake for dependency searching:

- Use a `PackageName_ROOT` variable to provide the location of a specific package.
- Add an install prefix to the `CMAKE_PREFIX_PATH` which is searched for all packages.
- Both the `PackageName_ROOT` and `CMAKE_PREFIX_PATH` can be used as either environment variables or CMake variables (passed via `-D`), where the CMake variable takes precedence.

```
# Several dependencies are installed under /opt/foo/bar and then a
# single dependency (HDF5 in this case) is installed in /opt/hdf5/1.13.0
$ export CMAKE_PREFIX_PATH=/opt/foo/bar
$ cmake -DHDF5_ROOT=/opt/hdf5/1.13.0 ../ADIOS2
```

Example: the following configuration will build, test and install under `/opt/adios2/2.9.0` an optimized (Release) version of ADIOS2.

```
$ cd build
$ cmake -DADIOS2_USE_Fortran=ON -DCMAKE_INSTALL_PREFIX=/opt/adios2/2.9.0 -DCMAKE_BUILD_
↪Type=Release ../ADIOS2
$ make -j16
$ ctest
$ make install
```

For a fully configurable build script, click [here](#).

## 4.3 Building on HPC Systems

1. **Modules:** Make sure all “module” dependencies are loaded and that minimum requirements are satisfied. Load the latest CMake module as many HPC systems default to an outdated version. Build with a C++11-compliant compiler, such as gcc >= 4.8.1, Intel >= 15, and PGI >= 15.
2. **Static/Dynamic build:** On Cray systems such as [Titan](#), the default behavior is static linkage, thus CMake builds ADIOS2 creates the static library `libadios2.a` by default. Read the system documentation to enable dynamic compilation, usually by setting an environment variable such as `CRAYPE_LINK_TYPE=dynamic`. Click [here](#) for a fully configurable script example on OLCF systems.
3. **Big Endian and 32-bit systems:** ADIOS2 hasn’t been tested on big endian and generally will not build on 32-bit systems. Please be aware before attempting to run.

4. **PGI compilers and C++11 support:** Version 15 of the PGI compiler is C++11 compliant. However it relies on the C++ standard library headers supplied by the system version of GCC, which may or may support all the C++11 features used in ADIOS2. On many systems (Titan at OLCF, for example) even though the PGI compiler supports C++11, the configured GCC and its headers do not (4.3.x on Cray Linux Environment, and v5 systems like Titan). To configure the PGI compiler to use a newer GCC, you must create a configuration file in your home directory that overrides the PGI compiler's default configuration. On Titan, the following steps will re-configure the PGI compiler to use GCC 6.3.0 provided by a module:

```
$ module load gcc/6.3.0
$ makelocalrc $(dirname $(which pgc++)) -gcc $(which gcc) -gpp $(which g++) -g77 $(which_
↪ gfortran) -o -net 1>${HOME}/.mypgirc 2>/dev/null
```

1. **Enabling RDMA for SST data transfers:** The SST engine in ADIOS2 is capable of using RDMA networks for transferring data between writer and reader cohorts, and generally this is the most performant data transport. However, SST depends upon libfabric to provide a generic interface to the underlying RDMA capabilities of the network, and properly configuring libfabric can be a difficult and error-prone task. HPC computing resources tend to be one-off custom resources with their own idiosyncracies, so this documentation cannot offer a definitive guide for every situation, but we can provide some general guidance and some recommendations for specific machines. If you are unable to configure ADIOS2 and libfabric to use RDMA, the best way to get help is to open an issue on the ADIOS2 github repository.

Pre-build concerns of note:

- on some HPC resources, libfabric is available as a loadable module. That should not be taken as an indication that that build of libfabric will work with SST, or even that it is compatible with the system upon which you find it. Your mileage may vary and you may have to build libfabric manually.
- libfabric itself depends upon other libraries like libibverbs and librdmacm. If you build libfabric with a package manager like spack, spack may build custom versions of those libraries as well, which may conflict with the system versions of those libraries.
- MPI on your HPC system may use libfabric itself, and linking your application with a different version of libfabric (or its dependent libraries) may result failure, possibly including opaque error messages from MPI.
- libfabric is structured in such a way that even if it is found during configuration, ADIOS *cannot* determine at compile time what providers will be present at run-time, or what their capabilities are. Therefore even a build that seems to successfully include libfabric and RDMA may be rejected at runtime as unable to support SST data transfer.

#### Configuration:

ADIOS2 uses the CMake `find_package()` functionality to locate libfabric. CMake will automatically search system libraries, but if you need to specify a libfabric location other than in a default system location you can add a “`-DLIBFABRIC_ROOT=<directory>`” argument to direct CMake to libfabric's location. If CMake finds libfabric, you should see the line “RDMA Transport for Staging: Available” near the end of the CMake output. This makes the RDMA DataTransport the default for SST data movement. (More information about SST engine parameters like *DataTransport* appears in the SST engine description.) If instead you see “RDMA Transport for Staging: Unconfigured”, RDMA will not be available to SST.

#### Run-time:

Generally, if RDMA is configured and the libfabric provider has the capabilities that SST needs for RDMA data transfer, SST will use RDMA without external acknowledgement. However, if RDMA is configured, but the libfabric provider doesn't have the capabilities that SST needs, ADIOS will output an error : ‘Warning: Preferred DataPlane “RDMA” not found.’ If you see this warning in a situation where you expect RDMA to be used, enabling verbose debugging output from SST may provide more information. The `SstVerbose` environment variable can have values from 1 to 5, with 1 being minimal debugging info (such as confirming which DataTransport is being used), and 5 being the most detailed debugging information from all ranks.

## 4.4 Installing the ADIOS2 library and the C++ and C bindings

By default, ADIOS2 will build the C++11 libadios2 library and the C and C++ bindings.

### 1. Minimum requirements:

- A C++11 compliant compiler
- An MPI C implementation on the syspath, or in a location identifiable by CMake.

### 2. Linking make install will copy the required headers and libraries into the directory specified by CMAKE\_INSTALL\_PREFIX:

- Libraries:
  - lib/libadios2.\* C++11 and C bindings
- Headers:
  - include/adios2.h C++11 namespace adios2
  - include/adios2\_c.h C prefix adios2\_
- Config file: run this command to get installation info
  - bin/adios2-config

## 4.5 Enabling the Python bindings

To enable the Python bindings in ADIOS2, based on [PyBind11](#), make sure to follow these guidelines:

### • Minimum requirements:

- Python 2.7 and above.
- *numpy*
- *mpi4py*

### • Running: If CMake enables Python compilation, an adios2.so library containing the Python module is generated in the build directory under lib/pythonX.X/site-packages/

- make sure your PYTHONPATH environment variable contains the path to adios2.so.
- make sure the Python interpreter is compatible with the version used for compilation via `python --version`.
- Run the Python tests with `ctest -R Python`
- Run [helloBPWriter.py](#) and [helloBPTimeWriter.py](#) via

```
$ mpirun -n 4 python helloBPWriter.py
$ python helloBPWriter.py
```

## 4.6 Enabling the Fortran bindings

### 1. Minimum requirements:

- A Fortran 90 compliant compiler
- A Fortran MPI implementation

### 2. Linking the Fortran bindings: `make install` will copy the required library and modules into the directory specified by `CMAKE_INSTALL_PREFIX`

- Library (note that `libadios2` must also be linked) - `lib/libadios2_f.*` - `lib/libadios2.*`
- Modules - `include/adios2/fortran/*.mod`

### 3. Module `adios2`: only module required to be used in an application use `adios`

## 4.7 Running Tests

ADIOS2 uses `googletest` to enable automatic testing after a CMake build. To run tests just type after building with `make`, run:

```
$ ctest
or
$ make test
```

The following screen will appear providing information on the status of each finalized test:

```
Test project /home/wfg/workspace/build
Start 1: ADIOSInterfaceWriteTest.DefineVarChar1x10
1/46 Test #1: ADIOSInterfaceWriteTest.DefineVarChar1x10 ..... Passed 0.06 sec
Start 2: ADIOSInterfaceWriteTest.DefineVarShort1x10
2/46 Test #2: ADIOSInterfaceWriteTest.DefineVarShort1x10 ..... Passed 0.04 sec
Start 3: ADIOSInterfaceWriteTest.DefineVarInt1x10
3/46 Test #3: ADIOSInterfaceWriteTest.DefineVarInt1x10 ..... Passed 0.04 sec
Start 4: ADIOSInterfaceWriteTest.DefineVarLong1x10
...
128/130 Test #128: ADIOSZfpWrapper.UnsupportedCall ..... Passed 0.05 sec
Start 129: ADIOSZfpWrapper.MissingMandatoryParameter
129/130 Test #129: ADIOSZfpWrapper.MissingMandatoryParameter ..... Passed 0.05 sec
Start 130: */TestManyVars.DontRedefineVars/*
130/130 Test #130: */TestManyVars.DontRedefineVars/* ..... Passed 0.08 sec

100% tests passed, 0 tests failed out of 130

Total Test time (real) = 204.82 sec
```

## 4.8 Running Examples

ADIOS2 is best learned by [examples](#).

A few very basic examples are described below:

Directory	Description
ADIOS2/examples/hello	very basic “hello world”-style examples for reading and writing <i>.bp</i> files.
ADIOS2/examples/ heatTransfer	2D Poisson solver for transients in Fourier’s model of heat transfer. Outputs <i>bp</i> , <i>dir</i> or <i>HDF5</i> .
ADIOS2/examples/basics	covers different <i>Variable</i> use cases classified by the dimension.

## AS PACKAGE

### 5.1 Conda

ADIOS2 can be obtained from anaconda cloud: [conda-forge adios2](#)

### 5.2 PyPI

ADIOS2 pip package can be downloaded with *pip3 install adios2* or *python3 -m pip install adios2*. This is contains the serial build only, so MPI programs cannot use it. See [adios2 on PyPi](#)

### 5.3 Spack

ADIOS2 is packaged in Spack. See [adios2 spack package](#)

### 5.4 Docker

Docker images including building and installation of dependencies and ADIOS 2 containers for Ubuntu 20 and CentOS 7 can be found in: under the directory [scripts/docker/](#)





## LINKING ADIOS 2

### 6.1 From CMake

ADIOS exports a CMake package configuration file that allows its targets to be directly imported into another CMake project via the `find_package` command:

```
cmake_minimum_required(VERSION 3.12)
project(MySimulation C CXX)

find_package(MPI REQUIRED)
find_package(ADIOS2 REQUIRED)
#...
add_library(my_library src1.cxx src2.cxx)
target_link_libraries(my_library PRIVATE adios2::cxx11_mpi MPI::MPI_C)
```

When configuring your project you can then set the `ADIOS2_ROOT` or `ADIOS2_DIR` environment variables to the install prefix of ADIOS2.

### 6.2 From non-CMake build systems

If you're not using CMake then you can manually get the necessary compile and link flags for your project using `adios2-config`:

```
$ /path/to/install-prefix/bin/adios2-config --cxxflags
ADIOS2_DIR: /path/to/install-prefix
-isystem /path/to/install-prefix/include -isystem /opt/ohpc/pub/mpi/openmpi3-gnu7/3.1.0/
↪include -pthread -std=gnu++11
$ /path/to/install-prefix/bin/adios2-config --cxxlibs
ADIOS2_DIR: /path/to/install-prefix
-Wl,-rpath,/path/to/install-prefix/lib:/opt/ohpc/pub/mpi/openmpi3-gnu7/3.1.0/lib /path/
↪to/install-prefix/lib/libadios2.so.2.4.0 -pthread -Wl,-rpath -Wl,/opt/ohpc/pub/mpi/
↪openmpi3-gnu7/3.1.0/lib -Wl,--enable-new-dtags -pthread /opt/ohpc/pub/mpi/openmpi3-
↪gnu7/3.1.0/lib/libmpi.so -Wl,-rpath-link,/path/to/install-prefix/lib
```



## USE ON DOE MACHINES

ADIOS2 is installed as part of the [E4S](#) software stack and access to adios2 is the same as access to the many other packages.

### 7.1 NERSC Perlmutter

To use adios2 on Perlmutter,

- load the e4s module
- pick your compiler environment with spack
- load adios2 with spack

```
~> module load e4s
```

---

The Extreme-Scale Scientific Software Stack (E4S) is accessible via the Spack package manager.

In order to access the production stack, you will need to load a spack environment. Here are some tips to get started:

```
'spack env list' - List all Spack environments
'spack env activate gcc' - Activate the "gcc" Spack environment
'spack env status' - Display the active Spack environment
'spack load amrex' - Load the "amrex" Spack package into your user environment
```

For additional support, please refer to the following references:

```
NERSC E4S Documentation: https://docs.nersc.gov/applications/e4s/
E4S Documentation: https://e4s.readthedocs.io
Spack Documentation: https://spack.readthedocs.io/en/latest/
Spack Slack: https://spackpm.slack.com
```

---

```
~> spack env list
==> 4 environments
   cce  cuda  gcc  nvhpc
~> spack env activate gcc
```

(continues on next page)

(continued from previous page)

```

~> spack load adios2

~> which bpls
/global/common/software/spackecp/perlmutter/e4s-23.08/94543/spack/opt/spack/linux-sles15-
↪zen3/gcc-12.3.0/adios2-2.9.1-iwv5lkkc5gyagr4uqrqr4v2fds7x66pk/bin/bpls

~> bpls -Vv
bpls: ADIOS file introspection utility

Build configuration:
ADIOS version: 2.9.1
C++ Compiler:  GNU 12.3.0 (CrayPrgEnv)
Target OS:      Linux-5.14.21-150400.24.81_12.0.87-cray_shasta_c
Target Arch:    x86_64
Available engines = 10: BP3, BP4, BP5, SST, SSC, Inline, MHS,
ParaViewADIOSInSituEngine, Null, Skeleton
Available operators = 4: BZip2, SZ, ZFP, PNG
Available features = 16: BP5, DATAMAN, MHS, SST, FORTRAN, MPI, BZIP2, PNG,
SZ, ZFP, O_DIRECT, CATALYST, SYSVSHMEM, ZEROMQ, PROFILING, ENDIAN_REVERSE

```

## 7.2 OLCF Frontier

OLCF installs the E4S packages in individual modules, hence *adios2* is also available as a module.

```

$ module avail adios2
----- /sw/frontier/spack-envs/base/modules/spack/cray-sles15-x86_64/cray-mpich/8.1.23-
↪j56azw5/cce/15.0.0 -----
adios2/2.8.1    adios2/2.8.3 (D)

Where:
D:  Default Module

$ module load adios2
$ bpls -Vv
bpls: ADIOS file introspection utility

Build configuration:
ADIOS version: 2.8.3
C++ Compiler:  GNU 12.2.0 (CrayPrgEnv)
Target OS:      Linux-5.14.21-150400.24.11_12.0.57-cray_shasta_c
Target Arch:    x86_64

```

## 7.3 ALCF Aurora

To use adios2 on Aurora,

- Load the default oneAPI (loaded automatically on login)
- `module use /soft/modulefiles`
- `module load spack-pe-oneapi/0.5-rc1`

This is a “metamodule” that makes many software packages from E4S loadable as modules.

```
$ module use /soft/modulefiles
$ module load spack-pe-oneapi/0.5-rc1
$ module avail adios2

----- /soft/packaging/spack/oneapi/0.5-rc1/modulefiles/Core -----
adios2/2.9.0-oneapi-mpich-testing
```



## INTERFACE COMPONENTS

### 8.1 Components Overview

---

**Note:** If you are doing simple tasks where performance is a non-critical aspect please go to the [High-Level APIs](#) section for a quick start. If you are an HPC application developer or you want to use ADIOS2 functionality in full please read this chapter.

---

The simple way to understand the big picture for the ADIOS2 unified user interface components is to map each class to the actual definition of the ADIOS acronym.

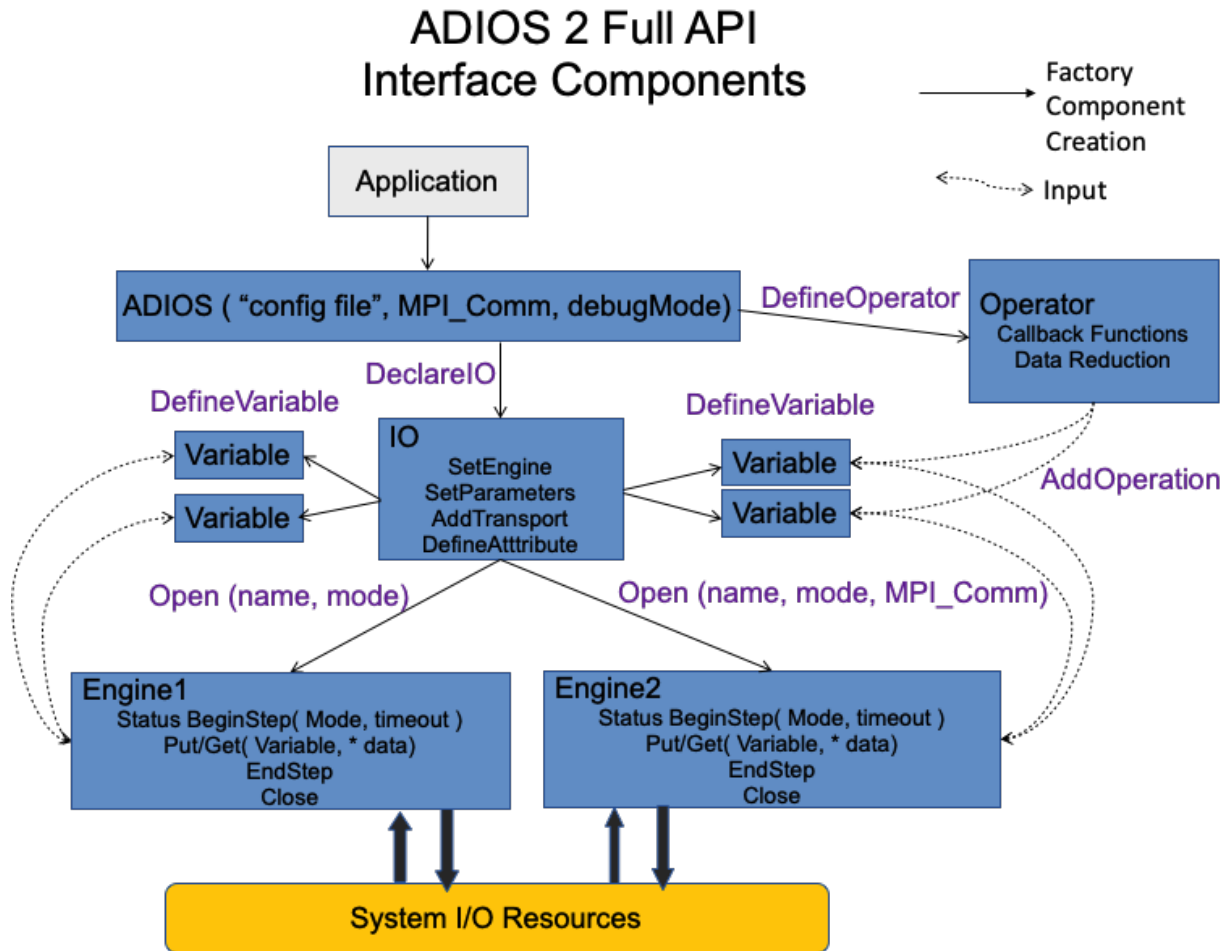
Component	Acronym	Function
<b>ADIOS</b>	ADaptable	Set MPI comm domain Set runtime settings Own other components
<b>IO</b>	I/O	Set engine Set variables/attributes Set compile-time settings
<b>Engine</b>	System	Execute heavy IO tasks Manage system resources

ADIOS2's public APIs are based on the natural choice for each supported language to represent each ADIOS2 components and its interaction with application datatypes. Thus,

Language	Component API	Application Data
C++(11/newer)	objects/member functions	pointers/references/std::vector
C	handler/functions	pointers
Fortran	handler/subroutines	arrays up to 6D
Python	objects/member functions	numpy arrays.

The following section provides a common overview to all languages based on the C++11 APIs. For each specific language go to the [Full APIs](#) section, but it's highly recommended to read this section as components map 1-to-1 in other languages.

The following figure depicts the components hierarchy from the application's point of view.



- **ADIOS:** the ADIOS component is the starting point between an application and the ADIOS2 library. Applications provide:

1. the scope of the ADIOS object through the MPI communicator,
2. an optional runtime configuration file (in XML format) to allow changing settings without recompiling.

The ADIOS component serves as a factory of adaptable IO components. Each IO must have a unique name within the scope of the ADIOS class object that created them with the DeclareIO function.

- **IO:** the IO component is the bridge between the application specific settings, transports. It also serves as a factory of:

1. Variables
2. Attributes
3. Engines

- **Variable:** Variables are the link between self-describing representation in the ADIOS2 library and data from applications. Variables are identified by unique names in the scope of the particular IO that created them. When the Engine API functions are called, a Variable must be provided along with the application data.

- **Attribute:** Attributes add extra information to the overall variables dataset defined in the IO class. They can be single or array values.

- **Engine:** Engines define the actual system executing the heavy IO tasks at Open, BeginStep, Put, Get, EndStep and Close. Due to polymorphism, new IO system solutions can be developed quickly reusing internal components



and reusing the same API. If `IO.SetEngine` is not called, the default engine is the binary-pack bp file reader and writer: **BPFile**.

- **Operator:** These define possible operations to be applied on adios2-managed data, for example, compression. This higher level abstraction is needed to provide support for callbacks, transforms, analytics, data models, etc. Any required task will be executed within the Engine. One or many operators can be associated with any of the adios2 objects or a group of them.

## 8.2 ADIOS

The `adios2::ADIOS` component is the initial contact point between an application and the ADIOS2 library. Applications can be classified as MPI and non-MPI based. We start by focusing on MPI applications as their non-MPI equivalent just removes the MPI communicator.

```
/** ADIOS class factory of IO class objects */
adios2::ADIOS adios("config.xml", MPI_COMM_WORLD);
```

This component is created by passing :

1. **Runtime config file** (optional): ADIOS2 xml runtime config file, see [Runtime Configuration Files](#).
2. **MPI communicator** : which determines the scope of the ADIOS library components in an application.

`adios2::ADIOS` objects can be created in MPI and non-MPI (serial) mode. Optionally, a runtime configuration file can be passed to the constructor indicating the full file path, name and extension.

### Constructors for MPI applications

```
/** Constructors */

// version that accepts an optional runtime adios2 config file
adios2::ADIOS(const std::string configFile,
              MPI_COMM mpiComm = MPI_COMM_SELF);

adios2::ADIOS(MPI_COMM mpiComm = MPI_COMM_SELF);

/** Examples */
adios2::ADIOS adios(MPI_COMM_WORLD);
adios2::ADIOS adios("config.xml", MPI_COMM_WORLD);
```

### Constructors for non-MPI (serial) applications

```
/** Constructors */
adios2::ADIOS(const std::string configFile);

adios2::ADIOS();

/** Examples */
adios2::ADIOS adios("config.xml");
adios2::ADIOS adios; // Do not use () for empty constructor.
```

**Factory of IO components:** Multiple IO components (IO tasks) can be created from within the scope of an ADIOS object by calling the `DeclareIO` function:

```
/** Signature */
adios2::IO ADIOS::DeclareIO(const std::string ioName);

/** Examples */
adios2::IO bpWriter = adios.DeclareIO("BPWriter");
adios2::IO bpReader = adios.DeclareIO("BPReader");
```

This function returns a reference to an existing IO class object that lives inside the ADIOS object that created it. The `ioName` string must be unique; declaring two IO objects with the same name will throw an exception. IO names are used to identify IO components in the runtime configuration file, [Runtime Configuration Files](#).

As shown in the diagram below, each resulting IO object is self-managed and independent, thus providing an adaptable way to perform different kinds of I/O operations. Users must be careful not to create conflicts between system level unique I/O identifiers: file names, IP address and port, MPI Send/Receive message rank and tag, etc.

---

**Tip:** The ADIOS component is the only one whose memory is owned by the application. Thus applications must decide on its scope. Any other component of the ADIOS2 API refers to a component that lives inside the ADIOS component (e.g. IO, Operator) or indirectly in the IO component (Variable, Engine)

---

## 8.3 IO

The IO component is the connection between how applications set up their input/output options by selecting an Engine and its specific parameters, subscribing variables to data, and setting supported transport modes to a particular Engine. Think of IO as a control panel for all the user-defined parameters that applications would like to fine tune. None of the IO operations are heavyweight until the Open function that generates an Engine is called. Its API allows

- generation of Variable and Attribute components containing information about the data in the input output process
- setting Engine-specific parameters and adding supported modes of transport
- generation of Engine objects to execute the actual IO tasks.

---

**Note:** If two different engine types are needed (e.g. BPFile, SST), you must define two IO objects. Also, at reading always define separate IOs to avoid Variable name clashes.

---

### 8.3.1 Setting a Particular Engine and its Parameters

Engines execute the heavy operations in ADIOS2. Each IO may select a type of Engine through the SetEngine function. If SetEngine is not called, then the BPFile engine is used.

```
/** Signature */
void adios2::IO::SetEngine( const std::string engineType );

/** Example */
bpIO.SetEngine("BPFile");
```

Each Engine allows the user to fine tune execution of buffering and output tasks via parameters passed to the IO object. These parameters are then propagated to the Engine. For a list of parameters allowed by each engine see [Available Engines](#).

---

**Note:** `adios2::Params` is an alias to `std::map<std::string, std::string>` to pass parameters as key-value string pairs, which can be initialized with curly-brace initializer lists.

---

```

/** Signature */
/** Passing several parameters at once */
void SetParameters(const adios2::Params& parameters);
/** Passing one parameter key-value pair at a time */
void SetParameter(const std::string key, const std::string value);

/** Examples */
io.SetParameters( { {"Threads", "4"},
                   {"ProfilingUnits", "Milliseconds"},
                   {"MaxBufferSize", "2Gb"},
                   {"BufferGrowthFactor", "1.5" }
                   {"FlushStepsCount", "5" }
                 } );
io.SetParameter( "Threads", "4" );

```

### 8.3.2 Adding Supported Transports with Parameters

The `AddTransport` function allows the user to specify how data is moved through the system, *e.g.* RDMA, wide-area networks, or files. It returns an unsigned `int` handler for each transport that can be used with the `Engine::Close` function at different times. `AddTransport` must provide library specific settings that the low-level system library interface allows.

```

/** Signature */
unsigned int AddTransport( const std::string transportType,
                          const adios2::Params& parameters );

/** Examples */
const unsigned int file1 = io.AddTransport( "File",
                                           { {"Library", "fstream"},
                                           {"Name", "file1.bp" }
                                           } );

const unsigned int file2 = io.AddTransport( "File",
                                           { {"Library", "POSIX"},
                                           {"Name", "file2.bp" }
                                           } );

const unsigned int wan = io.AddTransport( "WAN",
                                           { {"Library", "Zmq"},
                                           {"IP", "127.0.0.1" },
                                           {"Port", "80"}
                                           } );

```

### 8.3.3 Defining, Inquiring and Removing Variables and Attributes

The template functions `DefineVariable<T>` allows subscribing to data into ADIOS2 by returning a reference to a `Variable` class object whose scope is the same as the IO object that created it. The user must provide a unique name, the dimensions: MPI global: shape, MPI local: start and offset, optionally a flag indicating that dimensions are known to be constant, and a data pointer if defined in the application. Note: data is not passed at this stage. This is done by the Engine functions `Put` and `Get` for Variables. See the [Variable](#) section for supported types and shapes.

**Tip:** `adios2::Dims` is an alias to `std::vector<std::size_t>`, while `adios2::ConstantDims` is an alias to `bool`. Use them for code clarity.

```
/** Signature */
adios2::Variable<T>
    DefineVariable<T>(const std::string name,
                    const adios2::Dims &shape = {}, // Shape of global object
                    const adios2::Dims &start = {}, // Where to begin writing
                    const adios2::Dims &count = {}, // Where to end writing
                    const bool constantDims = false);

/** Example */
/** global array of floats with constant dimensions */
adios2::Variable<float> varFloats =
    io.DefineVariable<float>("bpFloats",
                           {size * Nx},
                           {rank * Nx},
                           {Nx},
                           adios2::ConstantDims);
```

Attributes are extra-information associated with the current IO object. The function `DefineAttribute<T>` allows for defining single value and array attributes. Keep in mind that Attributes apply to all Engines created by the IO object and, unlike Variables which are passed to each Engine explicitly, their definition contains their actual data.

```
/** Signatures */

/** Single value */
adios2::Attribute<T> DefineAttribute(const std::string &name,
                                    const T &value);

/** Arrays */
adios2::Attribute<T> DefineAttribute(const std::string &name,
                                    const T *array,
                                    const size_t elements);
```

In situations in which a variable and attribute has been previously defined: 1) a variable/attribute reference goes out of scope, or 2) when reading from an incoming stream, the IO can inquire about the status of variables and attributes. If the inquired variable/attribute is not found, then the overloaded `bool()` operator of returns `false`.

```
/** Signature */
adios2::Variable<T> InquireVariable<T>(const std::string &name) noexcept;
adios2::Attribute<T> InquireAttribute<T>(const std::string &name) noexcept;

/** Example */
```

(continues on next page)

(continued from previous page)

```

adios2::Variable<float> varPressure = io.InquireVariable<float>("pressure");
if( varPressure ) // it exists
{
    ...
}

```

**Note:** `adios2::Variable` overloads operator `bool()` so that we can check for invalid states (e.g. variables haven't arrived in a stream, weren't previously defined, or weren't written in a file).

**Caution:** Since `InquireVariable` and `InquireAttribute` are template functions, both the name and type must match the data you are looking for.

### 8.3.4 Opening an Engine

The `IO::Open` function creates a new derived object of the abstract `Engine` class and returns a reference handler to the user. A particular `Engine` type is set to the current `IO` component with the `IO::SetEngine` function. Engine polymorphism is handled internally by the `IO` class, which allows subclassing future derived `Engine` types without changing the basic API.

`Engine` objects are created in various modes. The available modes are `adios2::Mode::Read`, `adios2::Mode::ReadRandomAccess`, `adios2::Mode::ReadFlattenSteps`, `adios2::Mode::Write`, `adios2::Mode::Append`, `adios2::Mode::Sync`, `adios2::Mode::Deferred`, and `adios2::Mode::Undefined`.

```

/** Signatures */
/** Provide a new MPI communicator other than from ADIOS->IO->Engine */
adios2::Engine adios2::IO::Open(const std::string &name,
                               const adios2::Mode mode,
                               MPI_Comm mpiComm );

/** Reuse the MPI communicator from ADIOS->IO->Engine \n or non-MPI serial mode */
adios2::Engine adios2::IO::Open(const std::string &name,
                               const adios2::Mode mode);

/** Examples */

/** Engine derived class, spawned to start Write operations */
adios2::Engine bpWriter = io.Open("myVector.bp", adios2::Mode::Write);

/** Engine derived class, spawned to start Read operations on rank 0 */
if( rank == 0 )
{
    adios2::Engine bpReader = io.Open("myVector.bp",
                                      adios2::Mode::Read,
                                      MPI_COMM_SELF);
}

```

**Caution:** Always pass `MPI_COMM_SELF` if an Engine lives in only one MPI process. `Open` and `Close` are collective operations.

## 8.4 Variable

An `adios2::Variable` is the link between a piece of data coming from an application and its metadata. This component handles all application variables classified by data type and shape.

Each `IO` holds a set of Variables, and each `Variable` is identified with a unique name. They are created using the reference from `IO::DefineVariable<T>` or retrieved using the pointer from `IO::InquireVariable<T>` functions in *IO*.

### 8.4.1 Data Types

Only primitive types are supported in ADIOS2. Fixed-width types from `<cinttypes>` and `<cstdint>` should be preferred when writing portable code. ADIOS2 maps primitive types to equivalent fixed-width types (e.g. `int` -> `int32_t`). In C++, acceptable types `T` in `Variable<T>` along with their preferred fix-width equivalent in 64-bit platforms are given below:

Data types Variables supported by ADIOS2 `Variable<T>`

```
std::string (only used for global and local values, not arrays)
char -> int8_t or uint8_t depending on compiler flags
signed char -> int8_t
unsigned char -> uint8_t
short -> int16_t
unsigned short -> uint16_t
int -> int32_t
unsigned int -> uint32_t
long int -> int32_t or int64_t (Linux)
long long int -> int64_t
unsigned long int -> uint32_t or uint64_t (Linux)
unsigned long long int -> uint64_t
float -> always 32-bit = 4 bytes
double -> always 64-bit = 8 bytes
long double -> platform dependent
std::complex<float> -> always 64-bit = 8 bytes = 2 * float
std::complex<double> -> always 128-bit = 16 bytes = 2 * double
```

---

**Tip:** It's recommended to be consistent when using types for portability. If data is defined as a fixed-width integer, define variables in ADIOS2 using a fixed-width type, e.g. for `int32_t` data types use `DefineVariable<int32_t>`.

---

---

**Note:** C, Fortran APIs: the enum and parameter `adios2_type_XXX` only provides fixed-width types.

---

---

**Note:** Python APIs: use the equivalent fixed-width types from `numpy`. If `dtype` is not specified, ADIOS2 handles `numpy` defaults just fine as long as primitive types are passed.

---

## 8.4.2 Shapes

ADIOS2 is designed for MPI applications. Thus different application data shapes must be supported depending on their scope within a particular MPI communicator. The shape is defined at creation from the IO object by providing the dimensions: shape, start, count in the `IO::DefineVariable<T>`. The supported shapes are described below.

1. **Global Single Value:** Only a name is required for their definition. These variables are helpful for storing global information, preferably managed by only one MPI process, that may or may not change over steps: *e.g.* total number of particles, collective norm, number of nodes/cells, etc.

```
if( rank == 0 )
{
    adios2::Variable<uint32_t> varNodes = io.DefineVariable<uint32_t>("Nodes");
    adios2::Variable<std::string> varFlag = io.DefineVariable<std::string>(
    ↪ "Nodes flag");
    // ...
    engine.Put( varNodes, nodes );
    engine.Put( varFlag, "increased" );
    // ...
}
```

**Note:** Variables of type `string` are defined just like global single values. Multidimensional strings are supported for fixed size strings through variables of type `char`.

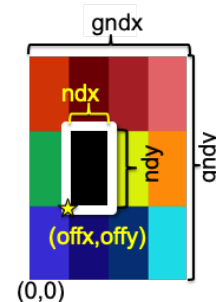
2. **Global Array:** This is the most common shape used for storing data that lives in several MPI processes. The image below illustrates the definitions of the dimension components in a global array: shape, start, and count.

### Variable Global

- N-dimensions
- Primitive Type
- Decomposition across many processors
  - global dimensions (Shape), local place (Start, Count)

```
adios2::Variable<double> &varT = io.DefineVariable<double>
(
    "T",           // name in output/input
    {gndx, gndy},  // Shape: physical dimensions (2D here)
    {offx, offy},  // Start: starting local offsets
    {ndx, ndy}     // Count: local size
);
```

- C/C++/Python always row-major, Fortran/Matlab/R always column-major



**Warning:** Be aware of data ordering in your language of choice (row-major or column-major) as depicted in the figure above. Data decomposition is done by the application, not by ADIOS2.

Start and Count local dimensions can be later modified with the `Variable::SetSelection` function if it is not a constant dimensions variable.

3. **Local Value:** Values that are local to the MPI process. They are defined by passing the `adios2::LocalValueDim` enum as follows:

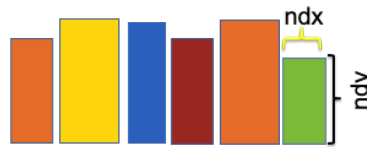
```
adios2::Variable<int32_t> varProcessID =
    io.DefineVariable<int32_t>("ProcessID", {adios2::LocalValueDim})
//...
engine.Put<int32_t>(varProcessID, rank);
```

These values become visible on the reader as a single merged 1-D Global Array whose size is determined by the number of writer ranks.

4. **Local Array:** Arrays that are local to the MPI process. These are commonly used to write checkpoint-restart data. Reading, however, needs to be handled differently: each process' array has to be read separately, using `SetSelection` per rank. The size of each process selection should be discovered by the reading application by inquiring per-block size information of the variable, and allocate memory accordingly.

## Variable Local

- N-dimensions
- Primitive Type
- Independent blocks on each rank
  - local dimensions (Count)



```
adios2::Variable<double> &varT = io.DefineVariable<double>
(
    "T",                // name in output/input
    {},
    {},
    {ndx, ndy}          // local size
);
```

- C/C++/Python always row-major, Fortran/Matlab/R always column-major

**Note:** Constants are not handled separately from step-varying values in ADIOS2. Simply write them only once from one rank.

5. **Joined Array:** Joined arrays are a variation of the Local Array described above. Where LocalArrays are only available to the reader via their block number, JoinedArrays are merged into a single global array whose global dimensions are determined by the sum of the contributions of each writer rank. Specifically: JoinedArrays are N-dimensional arrays where one (and only one) specific dimension is the Joined dimension. (The other dimensions must be constant and the same across all contributions.) When defining a Joined variable, one specifies a shape parameter that give the dimensionality of the array with the special constant `adios2::JoinedDim` in the dimension to be joined. Unlike a Global Array definition, the start parameter must be an empty Dims value. For example, the definition below defines a 2-D Joined array where the first dimension is the one along which blocks will be joined and the 2nd dimension is 5. Here this rank is contributing two rows to this array.

```
auto var = outIO.DefineVariable<double>("table", {adios2::JoinedDim, 5}, {}, {2, 5});
```

If each of N writer ranks were to declare a variable like this and do a single `Put()` in a timestep, the reader-side GlobalArray would have shape `{2*N, 5}` and all normal reader-side GlobalArray operations would be applicable to it.



---

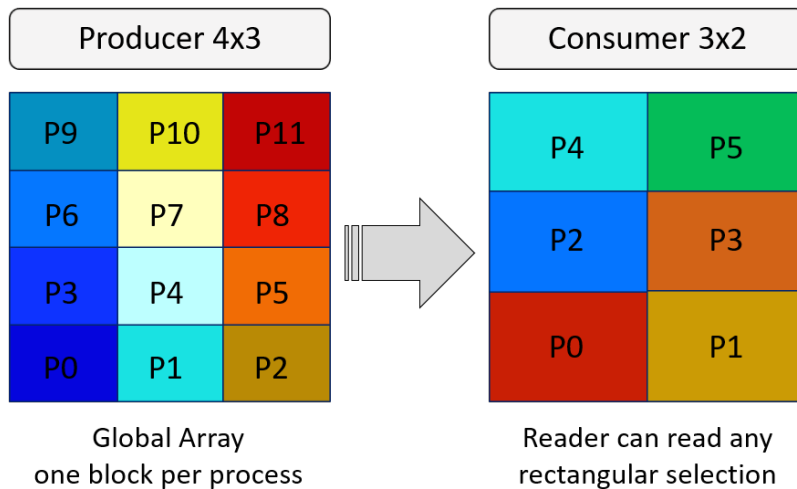
**Note:** JoinedArrays are currently only supported by BP4 and BP5 engines, as well as the SST engine with BP5 marshalling.

---

### 8.4.3 Global Array Capabilities and Limitations

ADIOS2 is focusing on writing and reading N-dimensional, distributed, global arrays of primitive types. The basic idea is that, usually, a simulation has such a data structure in memory (distributed across multiple processes) and wants to dump its content regularly as it progresses. ADIOS2 was designed to:

1. to do this writing and reading as fast as possible
2. to enable reading any subsection of the array



The figure above shows a parallel application of 12 processes producing a 2D array. Each process has a 2D array locally and the output is created by placing them into a 4x3 pattern. A reading application's individual process then can read any subsection of the entire global array. In the figure, a 6 process application decomposes the array in a 3x2 pattern and each process reads a 2D array whose content comes from multiple producer processes.

The figure hopefully helps to understand the basic concept but it can be also misleading if it suggests limitations that are not there. Global Array is simply a boundary in N-dimensional space where processes can place their blocks of data. In the global space:


1. one process can place multiple blocks

12 Producers  
multiple blocks per process

P9	P10	P11	P9	P7	P11
P6	P7	P8	P10	P11	P8
P3	P4	P5	P2	P4	P5
P0	P1	P2	P9	P5	P6

2. does NOT need to be fully covered by the blocks

Global Array sparsely filled out




“White” space is not stored in ADIOS2.  
Read returns 0 for those cells.

- at reading, unfilled positions will not change the allocated memory

3. blocks can overlap

## Global Array with overlapping blocks

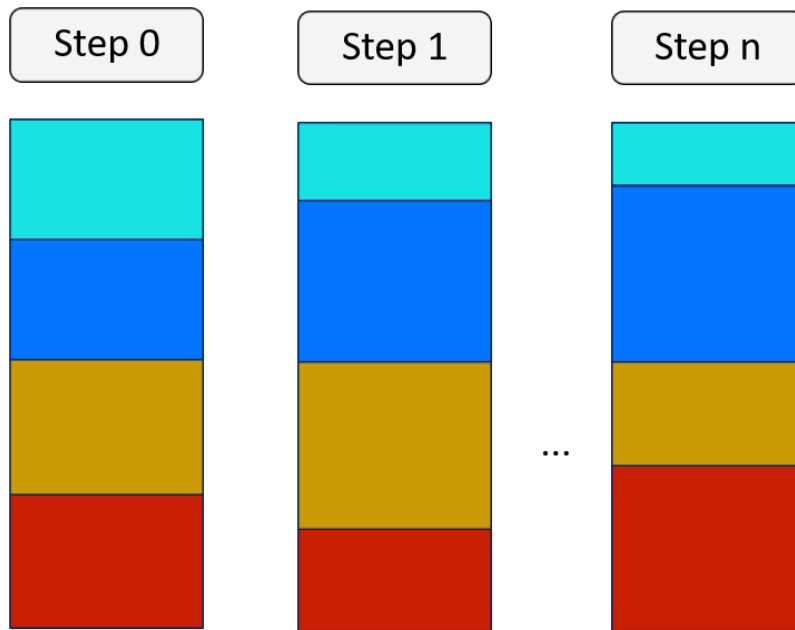


An overlapped cell has  
“one of the” values

- the reader will get values in an overlapping position from one of the block but there is no control over from which block
4. each process can put a different size of block, or put multiple blocks of different sizes
  5. some process may not contribute anything to the global array

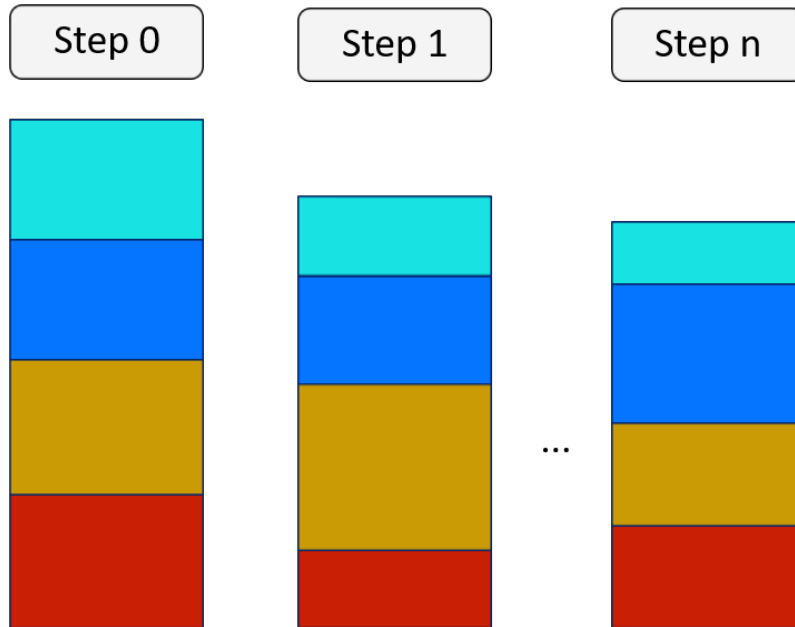
Over multiple output steps

1. the processes CAN change the size (and number) of blocks in the array
  - E.g. atom table: global size is fixed but atoms wander around processes, so their block size is changing



2D table where  
the number of rows per producer is changing

2. the global dimensions CAN change over output steps
  - but then you cannot read multiple steps at once
  - E.g. particle table size changes due to particles disappearing or appearing



2D table where  
the number of rows per producer as well as  
the total global size are changing

Limitations of the ADIOS2 global array concept

1. Indexing starts from 0
2. Cyclic data patterns are not supported; only blocks can be written or read
3. If Some blocks may fully or partially fall outside of the global boundary, the reader will not be able to read those parts

**Note:** Technically, the content of the individual blocks is kept in the BP format (but not in HDF5 format) and in staging. If you really, really want to retrieve all the blocks, you need to handle this array as a Local Array and read the blocks one by one.

## 8.5 Attribute

Attributes are extra information associated with a particular IO component. They can be thought of as a very simplified Variable, but with the goal of adding extra metadata. The most common use is the addition of human-readable metadata (e.g. "experiment name", "date and time", "04,27,2017", or a schema).

Currently, ADIOS2 supports single values and arrays of primitive types (excluding `complex<T>`) for the template type in the `IO::DefineAttribute<T>` and `IO::InquireAttribute<T>` function (in C++).

The data types supported for ADIOS2 Attributes are

```
std::string
char
```

(continues on next page)

(continued from previous page)

```
signed char
unsigned char
short
unsigned short
int
unsigned int
long int
long long int
unsigned long int
unsigned long long int
float
double
long double
```

The returned object (`DefineAttribute` or `InquireAttribute`) only serves the purpose to inspect the current `Attribute<T>` information within code.

## 8.6 Engine

The Engine abstraction component serves as the base interface to the actual IO systems executing the heavy-load tasks performed when producing and consuming data.

Engine functionality works around two concepts:

1. Variables are published (Put) and consumed (Get) in “steps” in either “File” random-access (all steps are available) or “Streaming” (steps are available as they are produced in a step-by-step fashion).
2. Variables are published (Put) and consumed (Get) using a “sync” or “deferred” (lazy evaluation) policy.

**Caution:** The ADIOS2 “step” is a logical abstraction that means different things depending on the application context. Examples: “time step”, “iteration step”, “inner loop step”, or “interpolation step”, “variable section”, etc. It only indicates how the variables were passed into ADIOS2 (e.g. I/O steps) without the user having to index this information on their own.

---

**Tip:** Publishing and consuming data is a round-trip in ADIOS2. Put and Get APIs for write/append and read modes aim to be “symmetric”, reusing functions, objects, and semantics as much as possible.

---

The rest of the section explains the important concepts.

### 8.6.1 BeginStep

Begins a logical step and return the status (via an enum) of the stream to be read/written. In streaming engines `BeginStep` is where the receiver tries to acquire a new step in the reading process. The full signature allows for a mode and timeout parameters. See *Supported Engines* for more information on what engine allows. A simplified signature allows each engine to pick reasonable defaults.

```
// Full signature
StepStatus BeginStep(const StepMode mode,
```

(continues on next page)

(continued from previous page)

```

const float timeoutSeconds = -1.f);

// Simplified signature
StepStatus BeginStep();

```

## 8.6.2 EndStep

Ends logical step, flush to transports depending on IO parameters and engine default behavior.

**Tip:** To write portable code for a step-by-step access across ADIOS2 engines (file and streaming engines) use `BeginStep` and `EndStep`.

**Danger:** Accessing random steps in read mode (e.g. `Variable<T>::SetStepSelection` in file engines) will create a conflict with `BeginStep` and `EndStep` and will throw an exception. In file engines, data is either consumed in a random-access or step-by-step mode, but not both.

## 8.6.3 Close

Close current engine and underlying transports. An `Engine` object can't be used after this call.

## 8.6.4 Put: modes and memory contracts

`Put` publishes data in ADIOS2. It is unavailable unless the `Engine` is created in `Write` or `Append` mode.

The most common signature is the one that passes a `Variable<T>` object for the metadata, a `const` piece of contiguous memory for the data, and a mode for either `Deferred` (data may be collected at `Put()` or not until `EndStep/PerformPuts/Close`) or `Sync` (data is reusable immediately). This is the most common use case in applications.

1. Deferred (default) or Sync mode, data is contiguous memory

```

void Put(Variable<T> variable, const T* data, const adios2::Mode =
↳ adios2::Mode::Deferred);

```

ADIOS2 Engines also provide direct access to their buffer memory. `Variable<T>::Span` is based on a subset of the upcoming C++20 `std::span`, which is a non-owning reference to a block of contiguous memory. Spans act as a 1D container meant to be filled out by the application. They provide the standard API of an STL container, providing `begin()` and `end()` iterators, `operator[]` and `at()`, as well as `data()` and `size()`.

`Variable<T>::Span` is helpful in situations in which temporaries are needed to create contiguous pieces of memory from non-contiguous pieces (e.g. tables, arrays without ghost-cells), or just to save memory as the returned `Variable<T>::Span` can be used for computation, thus avoiding an extra copy from user memory into the ADIOS2 buffer. `Variable<T>::Span` combines a hybrid `Sync` and `Deferred` mode, in which the initial value and memory allocations are `Sync`, while data population and metadata collection are done at `EndStep/PerformPuts/Close`. Memory contracts are explained later in this chapter followed by examples.

The following `Variable<T>::Span` signatures are available:

2. Return a span setting a default `T()` value into a default buffer

```
Variable<T>::Span Put(Variable<T> variable);
```

3. Return a span setting an initial fill value into a certain buffer. If span is not returned then the `fillValue` is fixed for that block.

```
Variable<T>::Span Put(Variable<T> variable, const size_t bufferID, const T_
↪fillValue);
```

In summary, the following are the current Put signatures for publishing data in ADIOS 2:

1. Deferred (default) or Sync mode, data is contiguous memory put in an ADIOS2 buffer.

```
void Put(Variable<T> variable, const T* data, const adios2::Mode =_
↪adios2::Mode::Deferred);
```

2. Return a span setting a default `T()` value into a default ADIOS2 buffer. If span is not returned then the default `T()` is fixed for that block (e.g. zeros).

```
Variable<T>::Span Put(Variable<T> variable);
```

3. Return a span setting an initial fill value into a certain buffer. If span is not returned then the `fillValue` is fixed for that block.

```
Variable<T>::Span Put(Variable<T> variable, const size_t bufferID, const T_
↪fillValue);
```

The following table summarizes the memory contracts required by ADIOS2 engines between Put signatures and the data memory coming from an application:

Put	Data Memory	Contract
De-ferred	Pointer	do not modify until PerformPuts/EndStep/Close
	Contents	consumed at Put or PerformPuts/EndStep/Close
Sync	Pointer	modify after Put
	Contents	consumed at Put
Span	Pointer	modified by new Spans, updated span iterators/data
	Contents	consumed at PerformPuts/EndStep/Close

**Note:** In Fortran (array) and Python (numpy array) avoid operations that modify the internal structure of an array (size) to preserve the address.

Each Engine will give a concrete meaning to each functions signatures, but all of them must follow the same memory contracts to the “data pointer”: the memory address itself, and the “data contents”: memory bits (values).

1. **Put in Deferred or lazy evaluation mode (default):** this is the preferred mode as it allows Put calls to be “grouped” before potential data transport at the first encounter of `PerformPuts`, `EndStep` or `Close`.

```
Put(variable, data);
Put(variable, data, adios2::Mode::Deferred);
```

Deferred memory contracts:

- “data pointer” do not modify (e.g. resize) until first call to `PerformPuts`, `EndStep` or `Close`.



- “data contents” may be consumed immediately or at first call to PerformPuts, EndStep or Close. Do not modify data contents after Put.

Usage:

```
// recommended use:
// set "data pointer" and "data contents"
// before Put
data[0] = 10;

// Puts data pointer into adios2 engine
// associated with current variable metadata
engine.Put(variable, data);

// Modifying data after Put(Deferred) may result in different
// results with different engines
// Any resize of data after Put(Deferred) may result in
// memory corruption or segmentation faults
data[1] = 10;

// "data contents" must not have been changed
// "data pointer" must be the same as in Put
engine.EndStep();
//engine.PerformPuts();
//engine.Close();

// now data pointer can be reused or modified
```

**Tip:** It’s recommended practice to set all data contents before Put in deferred mode to minimize the risk of modifying the data pointer (not just the contents) before PerformPuts/EndStep/Close.

2. **Put in Sync mode:** this is the special case, data pointer becomes reusable right after Put. Only use it if absolutely necessary (e.g. memory bound application or out of scope data, temporary).

```
Put(variable, *data, adios2::Mode::Sync);
```

Sync memory contracts:

- “data pointer” and “data contents” can be modified after this call.

Usage:

```
// set "data pointer" and "data contents"
// before Put in Sync mode
data[0] = 10;

// Puts data pointer into adios2 engine
// associated with current variable metadata
engine.Put(variable, data, adios2::Mode::Sync);

// data pointer and contents can be reused
// in application
```

3. **Put returning a Span:** signature that allows access to adios2 internal buffer.

Use cases:

- population from non-contiguous memory structures
- memory-bound applications

Limitations:

- does not allow operations (compression)
- must keep engine and variables within scope of span usage

Span memory contracts:

- “data pointer” provided by the engine and returned by `span.data()`, might change with the generation of a new span. It follows iterator invalidation rules from `std::vector`. Use `span.data()` or iterators, `span.begin()`, `span.end()` to keep an updated data pointer.
- span “data contents” are published at the first call to `PerformPuts`, `EndStep` or `Close`

Usage:

```
// return a span into a block of memory
// set memory to default T()
adios2::Variable<int32_t>::Span span1 = Put(var1);

// just like with std::vector::data()
// iterator invalidation rules
// dataPtr might become invalid
// always use span1.data() directly
T* dataPtr = span1.data();

// set memory value to -1 in buffer 0
adios2::Variable<float>::Span span2 = Put(var2, 0, -1);

// not returning a span just sets a constant value
Put(var3);
Put(var4, 0, 2);

// fill span1
span1[0] = 0;
span1[1] = 1;
span1[2] = 2;

// fill span2
span2[1] = 1;
span2[2] = 2;

// here collect all spans
// they become invalid
engine.EndStep();
//engine.PerformPuts();
//engine.Close();

// var1 = { 0, 1, 2 };
// var2 = { -1., 1., 2.};
// var3 = { 0, 0, 0};
// var4 = { 2, 2, 2};
```

The data fed to the Put function is assumed to be allocated on the Host (default mode). In order to use data allocated on the device, the memory space of the variable needs to be set to Cuda.

```
variable.SetMemorySpace(adios2::MemorySpace::CUDA);
engine.Put(variable, gpuData, mode);
```

**Note:** Only CUDA allocated buffers are supported for device data. Only the BP4 and BP5 engines are capable of receiving device allocated buffers.

### 8.6.5 PerformPuts

Executes all pending Put calls in deferred mode and collects span data. Specifically this call copies Put(Deferred) data into internal ADIOS buffers, as if Put(Sync) had been used instead.

**Note:** This call allows the reuse of user buffers, but may negatively impact performance on some engines.

### 8.6.6 PerformDataWrite

If supported by the engine, moves data from prior Put calls to disk

**Note:** Currently only supported by the BP5 file engine.

### 8.6.7 Get: modes and memory contracts

Get is the function for consuming data in ADIOS2. It is available when an Engine is created using Read mode at IO::Open. ADIOS2 Put and Get semantics are as symmetric as possible considering that they are opposite operations (*e.g.* Put passes `const T*`, while Get populates a non-const `T*`).

The Get signatures are described below.

1. Deferred (default) or Sync mode, data is contiguous pre-allocated memory:

```
Get(Variable<T> variable, const T* data, const adios2::Mode =
↪ adios2::Mode::Deferred);
```

2. In this signature, dataV is automatically resized by ADIOS2 based on the Variable selection:

```
Get(Variable<T> variable, std::vector<T>& dataV, const adios2::Mode =
↪ adios2::Mode::Deferred);
```

The following table summarizes the memory contracts required by ADIOS2 engines between Get signatures and the pre-allocated (except when using C++11 `std::vector`) data memory coming from an application:

Get	Data Memory	Contract
De-ferred	Pointer	do not modify until PerformGets/EndStep/Close
	Contents	populated at Get or PerformGets/EndStep/Close
Sync	Pointer	modify after Get
	Contents	populated at Get

1. **Get in Deferred or lazy evaluation mode (default):** this is the preferred mode as it allows Get calls to be “grouped” before potential data transport at the first encounter of PerformPuts, EndStep or Close.

```
Get(variable, data);
Get(variable, data, adios2::Mode::Deferred);
```

Deferred memory contracts:

- “data pointer”: do not modify (e.g. resize) until first call to PerformPuts, EndStep or Close.
- “data contents”: populated at Put, or at first call to PerformPuts, EndStep or Close.

Usage:

```
std::vector<double> data;

// resize memory to expected size
data.resize(varBlockSize);
// valid if all memory is populated
// data.reserve(varBlockSize);

// Gets data pointer to adios2 engine
// associated with current variable metadata
engine.Get(variable, data.data() );

// optionally pass data std::vector
// leave resize to adios2
//engine.Get(variable, data);

// "data pointer" must be the same as in Get
engine.EndStep();
// "data contents" are now ready
//engine.PerformPuts();
//engine.Close();

// now data pointer can be reused or modified
```

2. **Put in Sync mode:** this is the special case, data pointer becomes reusable right after Put. Only use it if absolutely necessary (e.g. memory bound application or out of scope data, temporary).

```
Get(variable, *data, adios2::Mode::Sync);
```

Sync memory contracts:

- “data pointer” and “data contents” can be modified after this call.

Usage:

```
.. code-block:: c++

std::vector<double> data;

// resize memory to expected size
data.resize(varBlockSize);
// valid if all memory is populated
// data.reserve(varBlockSize);
```

(continues on next page)

(continued from previous page)

```
// Gets data pointer to adios2 engine
// associated with current variable metadata
engine.Get(variable, data.data() );

// "data contents" are ready
// "data pointer" can be reused by the application
```

---

**Note:** Get doesn't support returning spans.

---

### 8.6.8 PerformGets

Executes all pending Get calls in deferred mode.

### 8.6.9 Engine usage example

The following example illustrates the basic API usage in write mode for data generated at each application step:

```
adios2::Engine engine = io.Open("file.bp", adios2::Mode::Write);

for( size_t i = 0; i < steps; ++i )
{
    // ... Application *data generation

    engine.BeginStep(); //next "logical" step for this application

    engine.Put(varT, dataT, adios2::Mode::Sync);
    // dataT memory already consumed by engine
    // Application can modify dataT address and contents

    // deferred functions return immediately (lazy evaluation),
    // dataU, dataV and dataW pointers and contents must not be modified
    // until PerformPuts, EndStep or Close.
    // 1st batch
    engine.Put(varU, dataU);
    engine.Put(varV, dataV);

    // in this case adios2::Mode::Deferred is redundant,
    // as this is the default option
    engine.Put(varW, dataW, adios2::Mode::Deferred);

    // effectively dataU, dataV, dataW are "deferred"
    // possibly until the first call to PerformPuts, EndStep or Close.
    // Application MUST NOT modify the data pointer (e.g. resize
    // memory) or change data contents.
    engine.PerformPuts();

    // dataU, dataV, dataW pointers/values can now be reused
```

(continues on next page)

(continued from previous page)

```

// ... Application modifies dataU, dataV, dataW

//2nd batch
dataU[0] = 10
dataV[0] = 10
dataW[0] = 10
engine.Put(varU, dataU);
engine.Put(varV, dataV);
engine.Put(varW, dataW);
// Application MUST NOT modify dataU, dataV and dataW pointers (e.g. resize),
// Contents should also not be modified after Put() and before
// PerformPuts() because ADIOS may access the data immediately
// or not until PerformPuts(), depending upon the engine
engine.PerformPuts();

// dataU, dataV, dataW pointers/values can now be reused

// Puts a varP block of zeros
adios2::Variable<double>::Span spanP = Put<double>(varP);

// Not recommended mixing static pointers,
// span follows
// the same pointer/iterator invalidation
// rules as std::vector
T* p = spanP.data();

// Puts a varMu block of 1e-6
adios2::Variable<double>::Span spanMu = Put<double>(varMu, 0, 1e-6);

// p might be invalidated
// by a new span, use spanP.data() again
foo(spanP.data());

// Puts a varRho block with a constant value of 1.225
Put<double>(varMu, 0, 1.225);

// it's preferable to start modifying spans
// after all of them are created
foo(spanP.data());
bar(spanMu.begin(), spanMu.end());

engine.EndStep();
// spanP, spanMu are consumed by the library
// end of current logical step,
// default behavior: transport data
}

engine.Close();
// engine is unreachable and all data should be transported
...

```

**Tip:** Prefer default `Deferred` (lazy evaluation) functions as they have the potential to group several variables with the trade-off of not being able to reuse the pointers memory space until `EndStep`, `PerformPuts`, `PerformGets`, or `Close`. Only use `Sync` if you really have to (*e.g.* reuse memory space from pointer). ADIOS2 prefers a step-based IO in which everything is known ahead of time when writing an entire step.

**Danger:** The default behavior of ADIOS2 `Put` and `Get` calls IS NOT synchronized, but rather deferred. It's actually the opposite of `MPI_Put` and more like `MPI_rPut`. Do not assume the data pointer is usable after a `Put` and `Get`, before `EndStep`, `Close` or the corresponding `PerformPuts/PerformGets`. Avoid using temporaries, r-values, and out-of-scope variables in `Deferred` mode. Use `adios2::Mode::Sync` in these cases.

## 8.6.10 Available Engines

A particular engine is set within the IO object that creates it with the `IO::SetEngine` function in a case insensitive manner. If the `SetEngine` function is not invoked the default engine is the `BPFile`.

Application	Engine	Description
File	BP5	DEFAULT write/read ADIOS2 native bp files
	HDF5	write/read interoperability with HDF5 files
Wide-Area-Network (WAN)	DataMan	write/read TCP/IP streams
Staging	SST	write/read to a “staging” area: <i>e.g.</i> RDMA

Engine polymorphism has two goals:

1. Each `Engine` implements an orthogonal IO scenario targeting a use case (*e.g.* Files, WAN, InSitu MPI, etc) using a simple, unified API.
2. Allow developers to build their own custom system solution based on their particular requirements in the own playground space. Reusable toolkit objects are available inside ADIOS2 for common tasks: bp buffering, transport management, transports, etc.

A class that extends `Engine` must be thought of as a solution to a range of IO applications. Each engine must provide a list of supported parameters, set in the IO object creating this engine using `IO::SetParameters`, and supported transports (and their parameters) in `IO::AddTransport`. Each Engine's particular options are documented in [Supported Engines](#).

## 8.7 Operator

The Operator abstraction allows ADIOS2 to act upon the user application data, either from a `adios2::Variable` or a set of `Variables` in an `adios2::IO` object. Current supported operations are:

1. Data compression/decompression, lossy and lossless.
2. Callback functions (C++11 bindings only) supported by specific engines

ADIOS2 enables the use of third-party libraries to execute these tasks.

Operators can be attached onto a variable in two modes: private or shared. In most situations, it is recommended to add an operator as a private one, which means it is owned by a certain variable. A simple example code is as follows.

```

#include <vector>
#include <adios2.h>
int main(int argc, char *argv[])
{
    std::vector<double> myData = {
        0.0001, 1.0001, 2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001,
        1.0001, 2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001,
        2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001,
        3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001,
        4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001,
        5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001,
        6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001,
        7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001,
        8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001, 1.0001,
        9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001, 1.0001, 0.0001,
    };
    adios2::ADIOS adios;
    auto io = adios.DeclareIO("TestIO");
    auto varDouble = io.DefineVariable<double>("varDouble", {10,10}, {0,0}, {10,10},
↪ adios2::ConstantDims);

    // add operator
    varDouble.AddOperation("mgard", {"accuracy", "0.01"});
    // end add operator

    auto engine = io.Open("hello.bp", adios2::Mode::Write);
    engine.Put<double>(varDouble, myData.data());
    engine.Close();
    return 0;
}

```

For users who need to attach a single operator onto multiple variables, a shared operator can be defined using the `adios2::ADIOS` object, and then attached to multiple variables using the reference to the operator object. Note that in this mode, all variables sharing this operator will also share the same configuration map. It should be only used when a number of variables need *exactly* the same operation. In real world use cases this is rarely seen, so please use this mode with caution.

```

#include <vector>
#include <adios2.h>
int main(int argc, char *argv[])
{
    std::vector<double> myData = {
        0.0001, 1.0001, 2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001,
        1.0001, 2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001,
        2.0001, 3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001,
        3.0001, 4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001,
        4.0001, 5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001,
        5.0001, 6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001,
        6.0001, 7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001,
        7.0001, 8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001,
        8.0001, 9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001, 1.0001,
        9.0001, 8.0001, 7.0001, 6.0001, 5.0001, 4.0001, 3.0001, 2.0001, 1.0001, 0.0001,
    };
}

```

(continues on next page)



(continued from previous page)

```

adios2::ADIOS adios;
auto io = adios.DeclareIO("TestIO");
auto varDouble = io.DefineVariable<double>("varDouble", {10,10}, {0,0}, {10,10},
↪ adios2::ConstantDims);

// define operator
auto op = adios.DefineOperator("SharedCompressor","mgard",{{"accuracy","0.01"}});
// add operator
varDouble.AddOperation(op);
// end add operator

auto engine = io.Open("hello.bp", adios2::Mode::Write);
engine.Put<double>(varDouble, myData.data());
engine.Close();
return 0;
}

```

**Warning:** Make sure your ADIOS2 library installation used for writing and reading was linked with a compatible version of a third-party dependency when working with operators. ADIOS2 will issue an exception if an operator library dependency is missing.

## 8.8 Runtime Configuration Files

ADIOS2 supports passing an optional runtime configuration file to the *ADIOS* component constructor (`adios2_init` in C, Fortran).

This file contains key-value pairs equivalent to the compile time `IO::SetParameters` (`adios2_set_parameter` in C, Fortran), and `IO::AddTransport` (`adios2_set_transport_parameter` in C, Fortran).

Each Engine and Operator must provide a set of available parameters as described in the *Supported Engines* section. Prior to version v2.6.0 only XML is supported; v2.6.0 and later support both XML and YAML.

**Warning:** Configuration files must have the corresponding format extension `.xml`, `.yaml`: `config.xml`, `config.yaml`, etc.

### 8.8.1 XML

```

<?xml version="1.0"?>
<adios-config>
  <io name="IONAME_1">

    <engine type="ENGINE_TYPE">

      <!-- Equivalent to IO::SetParameters-->
      <parameter key="KEY_1" value="VALUE_1"/>
      <parameter key="KEY_2" value="VALUE_2"/>
      <!-- ... -->
    
```

(continues on next page)

(continued from previous page)

```

    <parameter key="KEY_N" value="VALUE_N"/>

  </engine>

  <!-- Equivalent to IO::AddTransport -->
  <transport type="TRANSPORT_TYPE">
    <!-- Equivalent to IO::SetParameters-->
    <parameter key="KEY_1" value="VALUE_1"/>
    <parameter key="KEY_2" value="VALUE_2"/>
    <!-- ... -->
    <parameter key="KEY_N" value="VALUE_N"/>
  </transport>
</io>

<io name="IONAME_2">
  <!-- ... -->
</io>
</adios-config>

```

## 8.8.2 YAML

Starting with v2.6.0, ADIOS2 supports YAML configuration files. The syntax follows strict use of the YAML node keywords mapping to the ADIOS2 components hierarchy. If a keyword is unknown ADIOS2 simply ignores it. For an example file refer to `adios2 config file example` in our repo.

```

---
# adios2 config.yaml
# IO YAML Sequence (-) Nodes to allow for multiple IO nodes
# IO name referred in code with DeclareIO is mandatory

- IO: "IOName"

Engine:
  # If Type is missing or commented out, default Engine is picked up
  Type: "BP5"
  # optional engine parameters
  key1: value1
  key2: value2
  key3: value3

Variables:

  # Variable Name is Mandatory
  - Variable: "VariableName1"
    Operations:
      # Operation Type is mandatory (zfp, sz, etc.)
      - Type: operatorType
        key1: value1
        key2: value2

  - Variable: "VariableName2"

```

(continues on next page)

(continued from previous page)

```

Operations:
  # Operations sequence of maps
  - {Type: operatorType, key1: value1}
  - {Type: z-checker, key1: value1, key2: value2}

```

```

Transports:
  # Transport sequence of maps
  - {Type: file, Library: fstream}
  - {Type: rdma, Library: ibverbs}

```

```
...
```

**Caution:** YAML is case sensitive, make sure the node identifiers follow strictly the keywords: IO, Engine, Variables, Variable, Operations, Transports, Type.

**Tip:** Run a YAML validator or use a YAML editor to make sure the provided file is YAML compatible.

## 8.9 Anatomy of an ADIOS Program

### 8.9.1 Anatomy of an ADIOS Output

```

ADIOS adios("config.xml", MPI_COMM_WORLD);
|
|   IO io = adios.DeclareIO(...);
|   |
|   |   Variable<...> var = io.DefineVariable<...>(...)
|   |   Attribute<...> attr = io.DefineAttribute<...>(...)
|   |   Engine e = io.Open("OutputFileName.bp", adios2::Mode::Write);
|   |   |
|   |   |   e.BeginStep()
|   |   |   |
|   |   |   |   e.Put(var, datapointer);
|   |   |   |
|   |   |   e.EndStep()
|   |   |
|   |   e.Close();
|   |
|   |--> IO goes out of scope
|
|--> ADIOS goes out of scope or adios2_finalize()

```

The pseudo code above depicts the basic structure of performing output. The ADIOS object is necessary to hold all other objects. It is initialized with an MPI communicator in a parallel program or without in a serial program. Additionally, a config file (XML or YAML format) can be specified here to load runtime configuration. Only one ADIOS object is needed throughout the entire application but you can create as many as you want (e.g. if you need to separate IO objects using the same name in a program that reads similar input from an ensemble of multiple applications).

The IO object is required to hold the variable and attribute definitions, and runtime options for a particular input or output stream. The IO object has a name, which is used only to refer to runtime options in the configuration file. One IO object can only be used in one output or input stream. The only exception where an IO object can be used twice is one input stream plus one output stream where the output is reusing the variable definitions loaded during input.

Variable and Attribute definitions belong to one IO object, which means, they can only be used in one output. You need to define new ones for other outputs. Just because a Variable is defined, it will not appear in the output unless an associated Put() call provides the content.

A stream is opened and closed once. The Engine object implements the data movement for the stream. It depends on the runtime options of the IO object that what type of an engine is created in the Open() call. One output step is denoted by a pair of BeginStep..EndStep block.

An output step consist of variables and attributes. Variables are just definitions without content, so one must call a Put() function to provide the application data pointer that contains the data content one wants to write out. Attributes have their content in their definitions so there is no need for an extra call.

Some rules:

- Variables can be defined any time, before the corresponding Put() call
- Attributes can be defined any time before EndStep
- The following functions must be treated as Collective operations
- ADIOS
- Open
- BeginStep
- EndStep
- Close

---

**Note:** If there is only one output step, and we only want to write it to a file on disk, never stream it to other application, then BeginStep and EndStep are not required but it does not make any difference if they are called.

---

## 8.9.2 Anatomy of an ADIOS Input

```
ADIOS adios("config.xml", MPI_COMM_WORLD);
|
|   IO io = adios.DeclareIO(...);
|   |
|   |   Engine e = io.Open("InputFileName.bp", adios2::Mode::Read);
|   |   |
|   |   |   e.BeginStep()
|   |   |   |
|   |   |   |   varlist = io.AvailableVariables(...)
|   |   |   |   Variable var = io.InquireVariable(...)
|   |   |   |   Attribute attr = io.InquireAttribute(...)
|   |   |   |   |
|   |   |   |   |   e.Get(var, datapointer);
|   |   |   |   |
|   |   |   |   e.EndStep()
```

(continues on next page)

(continued from previous page)

```

| | |
| | | e.Close();
| | |
| | |--> IO goes out of scope
|
|--> ADIOS goes out of scope or adios2_finalize()

```

The difference between input and output is that while we have to define the variables and attributes for an output, we have to retrieve the available variables in an input first as definitions (Variable and Attribute objects).

If we know the particular variable (name and type) in the input stream, we can get the definition using `InquireVariable()`. Generic tools that process any input must use other functions to retrieve the list of variable names and their types first and then get the individual Variable objects. The same is true for Attributes.

### 8.9.3 Anatomy of an ADIOS File-only Input

Previously we explored how to read using the input mode `adios2::Mode::Read`. Nonetheless, ADIOS has another input mode named `adios2::Mode::ReadRandomAccess`. `adios2::Mode::Read` mode allows data access only timestep by timestep using `BeginStep/EndStep`, but generally it is more memory efficient as ADIOS is only required to load metadata for the current timestep. `ReadRandomAccess` can only be used with file engines and involves loading all the file metadata at once. So it can be more memory intensive than `adios2::Mode::Read` mode, but allows reading data from any timestep using `SetStepSelection()`. If you use `adios2::Mode::ReadRandomAccess` mode, be sure to allocate enough memory to hold multiple steps of the variable content. Note that ADIOS streaming engines (like SST, DataMan, etc.) do not support `ReadRandomAccess` mode. Also newer file Engines like BP5 to not allow `BeginStep/EndStep` calls in `ReadRandomAccess` mode.

```

ADIOS adios("config.xml", MPI_COMM_WORLD);
|
| IO io = adios.DeclareIO(...);
|
| | Engine e = io.Open("InputFileName.bp", adios2::Mode::ReadRandomAccess);
| | |
| | | Variable var = io.InquireVariable(...)
| | | | var.SetStepSelection()
| | | | e.Get(var, datapointer);
| | |
| | |
| | e.Close();
|
|--> IO goes out of scope
|
|--> ADIOS goes out of scope or adios2_finalize()

```

Previously we explored how to read using the input mode `adios2::Mode::Read`. Nonetheless, ADIOS has another input mode named `adios2::Mode::ReadRandomAccess`. `adios2::Mode::Read` mode allows data access only timestep by timestep using `BeginStep/EndStep`, but generally it is more memory efficient as ADIOS is only required to load metadata for the current timestep. `ReadRandomAccess` can only be used with file engines and involves loading all the file metadata at once. So it can be more memory intensive than `adios2::Mode::Read` mode, but allows reading data from any timestep using `SetStepSelection()`. If you use `adios2::Mode::ReadRandomAccess` mode, be sure to allocate enough memory to hold multiple steps of the variable content. Note that ADIOS streaming engines (like SST, DataMan, etc.) do not support `ReadRandomAccess` mode. Also newer file Engines like BP5 to not allow `BeginStep/EndStep` calls in `ReadRandomAccess` mode.

```
ADIOS adios("config.xml", MPI_COMM_WORLD);
|
|   IO io = adios.DeclareIO(...);
|   |
|   |   Engine e = io.Open("InputFileName.bp", adios2::Mode::ReadRandomAccess);
|   |   |
|   |   |   Variable var = io.InquireVariable(...)
|   |   |   |   var.SetStepSelection()
|   |   |   |   e.Get(var, datapointer);
|   |   |   |
|   |   |   e.Close();
|   |
|   |--> IO goes out of scope
|
|--> ADIOS goes out of scope or adios2_finalize()
```

In addition to the two read modes discussed above, ADIOS has another input mode named *adios2::Mode::ReadFlattenSteps*. This is a highly specialized mode built that is unlikely to be of general utility, but we describe it for completeness. In *ReadFlattenSteps* mode, ADIOS loads all the metadata in the file upon Open (just like *ReadRandomAccess* mode, but everything that was written appears that it was output on the same step, regardless of how many steps actually appear in the file. This affects the operation of many reader-side ADIOS functions, including *Steps()*, *BlocksInfo()*, *Get()*, etc.

## SUPPORTED VIRTUAL ENGINE NAMES

This section provides a description of the Virtual Engines that can be used to set up an actual Engine with specific parameters. These virtual names are used for beginner users to simplify the selection of an engine and its parameters. The following I/O uses cases are supported by virtual engine names:

1. **File:** File I/O (Default engine).

This sets up the I/O for files. If the file name passed in `Open()` ends with “.bp”, then the BP5 engine will be used starting in v2.9.0. If it ends with “.h5”, the HDF5 engine will be used. For old .bp files (BP version 3 format), the BP3 engine will be used for reading (v2.4.0 and below).

2. **FileStream:** Online processing via files.

This allows a Consumer to concurrently read the data while the Producer is writing new output steps into it. The Consumer will wait for the appearance of the file itself in `Open()` (for up to one hour) and wait for the appearance of new steps in the file (in `BeginStep()` up to the specified timeout in that function).

3. **InSituAnalysis:** Streaming data to another application.

This sets up ADIOS for transferring data from a Producer to a Consumer application. The Producer and Consumer are synchronized at `Open()`. The Consumer will receive every single output step from the Producer, therefore, the Producer will block on output if the Consumer is slow.

4. **InSituVisualization::** Streaming data to another application without waiting for consumption.

This sets up ADIOS for transferring data from a Producer to a Consumer without ever blocking the Producer. The Producer will throw away all output steps that are not immediately requested by a Consumer. It will also not wait for a Consumer to connect. This kind of streaming is great for an interactive visualization session where the user wants to see the most current state of the application.

5. **CodeCoupling::** Streaming data between two applications for code coupling.

Producer and Consumer are waiting for each other in `Open()` and every step must be consumed. Currently, this is the same as in situ analysis.

These virtual engine names are used to select a specific engine and its parameters. In practice, after selecting the virtual engine name, one can modify the settings by adding/overwriting parameters. Eventually, a seasoned user would use the actual Engine names and parameterize it for the specific run.

## 9.1 Virtual Engine Setups

These are the actual settings in ADIOS2 when a virtual engine is selected. The parameters below can be modified before the Open call.

1. **File**. Refer to the parameter settings for these engines of BP5, BP4, BP3 and HDF5 engines earlier in this section.
2. **FileStream**. The engine is BP5. The parameters are set to:

Key	Value Format	Default and Examples
OpenTimeoutSecs	float	<b>3600</b> (wait for up to an hour)
BeginStepPollingFrequencySecs	float	<b>1</b> (poll the file system with 1 second frequency)

3. **InSituAnalysis**. The engine is SST. The parameters are set to:

Key	Value Format	Default and Examples
RendezvousReaderCount	integer	<b>1</b> (Producer waits for the Consumer in Open)
QueueLimit	integer	<b>1</b> (only buffer one step)
QueueFullPolicy	string	<b>Block</b> (wait for the Consumer to get every step)
FirstTimestepPrecious	bool	false (SST default)
AlwaysProvideLatestTimestep	bool	false (SST default)

4. **InSituVisualization**. The engine is SST. The parameters are set to:

Key	Value Format	Default and Examples
RendezvousReaderCount	integer	<b>0</b> (Producer does NOT wait for Consumer in Open)
QueueLimit	integer	<b>3</b> (buffer first step + last two steps)
QueueFullPolicy	string	<b>Discard</b> (slow Consumer will miss out on steps)
FirstTimestepPrecious	bool	<b>true</b> (First step is kept around for late Consumers)
AlwaysProvideLatestTimestep	bool	false (SST default)

5. **Code Coupling**. The engine is SST. The parameters are set to:

Key	Value Format	Default and Examples
RendezvousReaderCount	integer	<b>1</b> (Producer waits for the Consumer in Open)
QueueLimit	integer	<b>1</b> (only buffer one step)
QueueFullPolicy	string	<b>Block</b> (wait for the Consumer to get every step)
FirstTimestepPrecious	bool	false (SST default)
AlwaysProvideLatestTimestep	bool	false (SST default)



## SUPPORTED ENGINES

This section provides a description of the *Available Engines* in ADIOS2 and their specific parameters to allow extra-control from the user. Parameters are passed in key-value pairs for:

1. Engine specific parameters
2. Engine supported transports and parameters

Parameters are passed at:

1. Compile time `IO::SetParameters` (`adios2_set_parameter` in C, Fortran)
2. Compile time `IO::AddTransport` (`adios2_set_transport_parameter` in C, Fortran)
3. *Runtime Configuration Files* in the *ADIOS* component.

### 10.1 BP5

The BP5 Engine writes and reads files in ADIOS2 native binary-pack (bp version 5) format. This was a new format for ADIOS 2.8, improving on the metadata operations and the memory consumption of the older BP4/BP3 formats. BP5 is the default file format as of ADIOS 2.9. As compared to the older format, BP5 provides three main advantages:

- **Lower memory** consumption. Deferred Puts will use user buffer for I/O wherever possible thus saving on a memory copy. Aggregation uses a fixed-size shared-memory segment on each compute node instead of using MPI to send data from one process to another. Memory consumption can get close to half of BP4 in some cases.
- **Faster metadata** management improves write/read performance where hundreds or more variables are added to the output.
- Improved functionality around **appending** many output steps into the same file. Better performance than writing new files each step. Restart can append to an existing series by truncating unwanted steps. Readers can filter out unwanted steps to only see and process a limited set of steps. Just like as in BP4, existing steps cannot be corrupted by appending new steps.

In 2.8 BP5 was a brand new file format and engine. It still does **NOT** support some functionality of BP4:

- **Burst buffer support** for writing data.

BP5 files have the following structure given a “name” string passed as the first argument of `IO::Open`:

```
io.SetEngine("BP5");  
adios2::Engine bpFile = io.Open("name", adios2::Mode::Write);
```

will generate:

```
% BP5 datasets are always a directory
name.bp/

% data and metadata files
name.bp/
    data.0
    data.1
    ...
    data.M
    md.0
    md.idx
    mmd.0
```

---

**Note:** BP5 file names are compatible with the Unix (/) and Windows (\\) file system naming convention for directories and files.

---

---

**Note:** BP5 has an `mmd.0` file in the directory, which BP4 does not have.

---

This engine allows the user to fine tune the buffering operations through the following optional parameters:

1. Streaming through file
  1. **OpenTimeoutSecs:** (Streaming mode) Reader may want to wait for the creation of the file in `io.Open()`. By default the `Open()` function returns with an error if file is not found.
  2. **BeginStepPollingFrequencySecs:** (Streaming mode) Reader can set how frequently to check the file (and file system) for new steps. Default is 1 seconds which may be stressful for the file system and unnecessary for the application.
2. Aggregation
  1. **AggregationType:** *TwoLevelShm*, *EveryoneWritesSerial* and *EveryoneWrites* are three aggregation strategies. See [Aggregation in BP5](#). The default is *TwoLevelShm*.
  2. **NumAggregators:** The number of processes that will ever write data directly to storage. The default is set to the number of compute nodes the application is running on (i.e. one process per compute node). *TwoLevelShm* will select a fixed number of processes *per compute-node* to get close to the intention of the user but does not guarantee the exact number of aggregators.
  3. **AggregatorRatio:** An alternative option to `NumAggregators` to pick every *n*th process as aggregator. The number of aggregators will be automatically kept to be within 1 and total number of processes no matter what bad number is supplied here. Moreover, *TwoLevelShm* will select an fixed number of processes *per compute-node* to get close to the intention of this ratio but does not guarantee the exact number of aggregators.
  4. **NumSubFiles:** The number of data files to write to in the `.bp/` directory. Only used by *TwoLevelShm* aggregator, where the number of files can be smaller then the number of aggregators. The default is set to *NumAggregators*.
  5. **StripeSize:** The data blocks of different processes are aligned to this size (default is 4096 bytes) in the files. Its purpose is to avoid multiple processes to write to the same file system block and potentially slow down the write.
  6. **MaxShmSize:** Upper limit for how much shared memory an aggregator process in *TwoLevelShm* can allocate. For optimum performance, this should be at least  $2xM + 1KB$  where *M* is the maximum size any

process writes in a single step. However, there is no point in allowing for more than 4GB. The default is 4GB.

### 3. Buffering

1. **BufferVType:** *chunk* or *malloc*, default is chunking. Chunking maintains the buffer as a list of memory blocks, either ADIOS-owned for sync-ed Puts and small Puts, and user-owned pointers of deferred Puts. Malloc maintains a single memory block and extends it (reallocates) whenever more data is buffered. Chunking incurs extra cost in I/O by having to write data in chunks (multiple write system calls), which can be helped by increasing *BufferChunkSize* and *MinDeferredSize*. Malloc incurs extra cost by reallocating memory whenever more data is buffered (by *Put()*), which can be helped by increasing *InitialBufferSize*.
2. **BufferChunkSize:** (for *chunk* buffer type) The size of each memory buffer chunk, default is 128MB but it is worth increasing up to 2147381248 (a bit less than 2GB) if possible for maximum write performance.
3. **MinDeferredSize:** (for *chunk* buffer type) Small user variables are always buffered, default is 4MB.
4. **InitialBufferSize:** (for *malloc* buffer type) initial memory provided for buffering (default and minimum is 16Kb). To avoid reallocations, it is worth increasing this size to the expected maximum total size of data any process would write in any step (not counting deferred Puts).
5. **GrowthFactor:** (for *malloc* buffer type) exponential growth factor for initial buffer > 1, default = 1.05.

### 4. Managing steps

1. **AppendAfterSteps:** BP5 enables overwriting some existing steps by opening in *adios2::Mode::Append* mode and specifying how many existing steps to keep. Default value is `MAX_INT`, so it always appends after the last step. -1 would achieve the same thing. If you have 10 steps in the file,
  - value 0 means starting from the beginning, truncating all existing data
  - value 1 means appending after the first step, so overwrite 2,3...10
  - value 10 means appending after all existing steps
  - value >10 means the same, append after all existing steps (gaps in steps are impossible)
  - -1 means appending after the last step, i.e. same as 10 or higher
  - -2 means removing the last step, i.e. starting from the 10th
  - -11 (and <-11) means truncating all existing data
2. **SelectSteps:** BP5 reading allows for only seeing selected steps. This is a string of space-separated list of range definitions in the form of “start:end:step”. Indexing starts from 0. If ‘end’ is ‘n’ or ‘N’, then it is an unlimited range expression. Range definitions are adding up. Note that in the reading functions, counting the steps is *always* 0 to *s-1* where *s* steps are presented, so even after applying this selection, the selected steps are presented as 0 to *s-1*. Examples:
  - “0 6 3 2” selects four steps indexed 0,2,3 and 6 (presented in reading as 0,1,2,3)
  - “1:5” selects 5 consecutive steps, skipping step 0, and starting from 1
  - “2:n” selects all steps from step 2
  - “0:n:2” selects every other steps from the beginning (0,2,4,6...)
  - “0:n:3 10:n:5” selects every third step from the beginning and additionally every fifth steps from step 10.

### 5. Asynchronous writing I/O

1. **AsyncOpen:** *true/false* Call the open function asynchronously. It decreases I/O overhead when creating lots of subfiles (*NumAggregators* is large) and one calls *io.Open()* well ahead of the first write step. Only implemented for writing. Default is *true*.

2. **AsyncWrite:** *true/false* Perform data writing operations asynchronously after *EndStep()*. Default is *false*. If the application calls *EnterComputationBlock()/ExitComputationBlock()* to indicate phases where no communication is happening, ADIOS2 will try to perform all data writing during those phases, otherwise it will write immediately and eagerly after *EndStep()*.
6. Direct I/O. Experimental, see discussion on [GitHub](#).
    1. **DirectIO:** Turn on `O_DIRECT` when using POSIX transport. Do not use this on parallel file systems.
    2. **DirectIOAlignOffset:** Alignment for file offsets. Default is 512 which is usually
    3. **DirectIOAlignBuffer:** Alignment for memory pointers. Default is to be same as *DirectIOAlignOffset*.
  7. Miscellaneous
    1. **StatsLevel:** 1 turns on *Min/Max* calculation for every variable, 0 turns this off. Default is 1. It has some cost to generate this metadata so it can be turned off if there is no need for this information.
    2. **MaxOpenFilesAtOnce:** Specify how many subfiles a process can keep open at once. Default is unlimited. If a dataset contains more subfiles than how many open file descriptors the system allows (see *ulimit -n*) then one can either try to raise that system limit (set it with *ulimit -n*), or set this parameter to force the reader to close some subfiles to stay within the limits.
    3. **Threads:** Read side: Specify how many threads one process can use to speed up reading. The default value is 0, to let the engine estimate the number of threads based on how many processes are running on the compute node and how many hardware threads are available on the compute node but it will use maximum 16 threads. Value 1 forces the engine to read everything within the main thread of the process. Other values specify the exact number of threads the engine can use. Although multithreaded reading works in a single *Get(adios2::Mode::Sync)* call if the read selection spans multiple data blocks in the file, the best parallelization is achieved by using deferred mode and reading everything in *PerformGets()/EndStep()*.

Key	Value Format	Default and Examples
OpenTimeoutSecs	float	<b>0</b> for <i>ReadRandomAccess</i> mode, <b>3600</b> for <i>Read</i> mode, <b>10.0</b> , <b>5</b>
BeginStepPollingFrequency-Secs	float	<b>1</b> , 10.0
AggregationType	string	<b>TwoLevelShm</b> , EveryoneWritesSerial, EveryoneWrites
NumAggregators	integer >= 1	<b>0 (one file per compute node)</b>
AggregatorRatio	integer >= 1	not used unless set
NumSubFiles	integer >= 1	<b>=NumAggregators</b> , only used when <i>Aggregation-Type=TwoLevelShm</i>
StripeSize	integer+units	<b>4KB</b>
MaxShmSize	integer+units	<b>4294762496</b>
BufferVType	string	<b>chunk</b> , malloc
BufferChunkSize	integer+units	<b>128MB</b> , worth increasing up to min(2GB, data-size/process/step)
MinDeferredSize	integer+units	<b>4MB</b>
InitialBufferSize	float+units >= 16Kb	<b>16Kb</b> , 10Mb, 0.5Gb
GrowthFactor	float > 1	<b>1.05</b> , 1.01, 1.5, 2
AppendAfterSteps	integer >= 0	<b>INT_MAX</b>
SelectSteps	string	“0 6 3 2”, “1:5”, “0:n:3 10:n:5”
AsyncOpen	string On/Off	<b>On</b> , Off, true, false
AsyncWrite	string On/Off	<b>Off</b> , On, true, false
DirectIO	string On/Off	<b>Off</b> , On, true, false
DirectIOAlignOffset	integer >= 0	<b>512</b>
DirectIOAlignBuffer	integer >= 0	set to DirectIOAlignOffset if unset
StatsLevel	integer, 0 or 1	<b>1</b> , 0
MaxOpenFilesAtOnce	integer >= 0	<b>UINT_MAX</b> , 1024, 1
Threads	integer >= 0	<b>0</b> , 1, 32

Only file transport types are supported. Optional parameters for `I0::AddTransport` or in runtime config file transport field:

#### Transport type: File

Key	Value Format	Default and Examples
Library	string	<b>POSIX</b> (UNIX), <b>FStream</b> (Windows), <b>stdio</b> , <b>IME</b>

The IME transport directly reads and writes files stored on DDN’s IME burst buffer using the IME native API. To use the IME transport, IME must be available on the target system and ADIOS2 needs to be configured with `ADIOS2_USE_IME`. By default, data written to the IME is automatically flushed to the parallel filesystem at every `EndStep()` call. You can disable this automatic flush by setting the transport parameter `SyncToPFS` to `OFF`.

## 10.2 BP4

The BP4 Engine writes and reads files in ADIOS2 native binary-pack (bp version 4) format. This was a new format for ADIOS 2.5 and improved on the metadata operations of the older BP3 format. Compared to the older format, BP4 provides three main advantages:

- Fast and safe **appending** of multiple output steps into the same file. Better performance than writing new files each step. Existing steps cannot be corrupted by appending new steps.
- **Streaming** through files (i.e. online processing). Consumer apps can read existing steps while the Producer is still writing new steps. Reader's loop can block (with timeout) and wait for new steps to arrive. Same reader code can read the entire data in post or in situ. No restrictions on the Producer.
- **Burst buffer support** for writing data. It can write the output to a local file system on each compute node and drain the data to the parallel file system in a separate asynchronous thread. Streaming read from the target file system are still supported when data goes through the burst buffer. Appending to an existing file on the target file system is NOT supported currently.

BP4 files have the following structure given a "name" string passed as the first argument of `IO::Open`:

```
io.SetEngine("BP4");
adios2::Engine bpFile = io.Open("name", adios2::Mode::Write);
```

will generate:

```
% BP4 datasets are always a directory
name.bp/

% data and metadata files
name.bp/
    data.0
    data.1
    ...
    data.M
    md.0
    md.idx
```

---

**Note:** BP4 file names are compatible with the Unix (/) and Windows (\\) file system naming convention for directories and files.

---

This engine allows the user to fine tune the buffering operations through the following optional parameters:

1. **Profile:** turns ON/OFF profiling information right after a run
2. **ProfileUnits:** set profile units according to the required measurement scale for intensive operations
3. **Threads:** number of threads provided from the application for buffering, use this for very large variables in data size
4. **InitialBufferSize:** initial memory provided for buffering (minimum is 16Kb)
5. **BufferGrowthFactor:** exponential growth factor for initial buffer > 1, default = 1.05.
6. **MaxBufferSize:** maximum allowable buffer size (must be larger than 16Kb). If too large adios2 will throw an exception.

7. **FlushStepsCount**: users can select how often to produce the more expensive collective metadata file in terms of steps: default is 1. Increase to reduce adios2 collective operations footprint, with the trade-off of reducing checkpoint frequency. Buffer size will increase until first steps count if `MaxBufferSize` is not set.
8. **NumAggregators** (or **SubStreams**): Users can select how many sub-files (`M`) are produced during a run, ranges between 1 and the number of mpi processes from `MPI_Size` (`N`), adios2 will internally aggregate data buffers (`N-to-M`) to output the required number of sub-files. Default is 0, which will let adios2 to group processes per shared-memory-access (i.e. one per compute node) and use one process per node as an aggregator. If `NumAggregators` is larger than the number of processes then it will be set to the number of processes.
9. **AggregatorRatio**: An alternative option to `NumAggregators` to pick every `Nth` process as aggregator. An integer divider of the number of processes is required, otherwise a runtime exception is thrown.
10. **OpenTimeoutSecs**: (Streaming mode) Reader may want to wait for the creation of the file in `io.Open()`. By default the `Open()` function returns with an error if file is not found.
11. **BeginStepPollingFrequencySecs**: (Streaming mode) Reader can set how frequently to check the file (and file system) for new steps. Default is 1 seconds which may be stressful for the file system and unnecessary for the application.
12. **StatsLevel**: Turn on/off calculating statistics for every variable (Min/Max). Default is On. It has some cost to generate this metadata so it can be turned off if there is no need for this information.
13. **StatsBlockSize**: Calculate Min/Max for a given size of each process output. Default is one Min/Max per writer. More fine-grained min/max can be useful for querying the data.
14. **NodeLocal** or **Node-Local**: For distributed file system. Every writer process must make sure the `.bp/` directory is created on the local file system. Required when writing to local disk/SSD/NVMe in a cluster. Note: the `BurstBuffer*` parameters are newer and should be used for using the local storage as temporary instead of this parameter.
15. **BurstBufferPath**: Redirect output file to another location and drain it to the original target location in an asynchronous thread. It requires to be able to launch one thread per aggregator (see `SubStreams`) on the system. This feature can be used on machines that have local NVMe/SSDs on each node to accelerate the output writing speed. On Summit at OLCF, use `"/mnt/bb/<username>"` for the path where `<username>` is your user account name. Temporary files on the accelerated storage will be automatically deleted after the application closes the output and ADIOS drains all data to the file system, unless draining is turned off (see the next parameter). Note: at this time, this feature cannot be used to append data to an existing dataset on the target system.
16. **BurstBufferDrain**: To write only to the accelerated storage but to not drain it to the target file system, set this flag to false. Data will NOT be deleted from the accelerated storage on close. By default, setting the `BurstBufferPath` will turn on draining.
17. **BurstBufferVerbose**: Verbose level 1 will cause each draining thread to print a one line report at the end (to standard output) about where it has spent its time and the number of bytes moved. Verbose level 2 will cause each thread to print a line for each draining operation (file creation, copy block, write block from memory, etc).
18. **StreamReader**: By default the BP4 engine parses all available metadata in `Open()`. An application may turn this flag on to parse a limited number of steps at once, and update metadata when those steps have been processed. If the flag is ON, reading only works in streaming mode (using `BeginStep/EndStep`); file reading mode will not work as there will be zero steps processed in `Open()`.

Key	Value Format	Default and Examples
Profile	string On/Off	<b>On</b> , Off
ProfileUnits	string	<b>Microseconds</b> , Milliseconds, Seconds, Minutes, Hours
Threads	integer > 1	<b>1</b> , 2, 3, 4, 16, 32, 64
InitialBufferSize	float+units 16Kb	<b>&gt;= 16Kb</b> , 10Mb, 0.5Gb
MaxBufferSize	float+units 16Kb	<b>&gt;= at EndStep</b> , 10Mb, 0.5Gb
BufferGrowthFactor	float > 1	<b>1.05</b> , 1.01, 1.5, 2
FlushStepsCount	integer > 1	<b>1</b> , 5, 1000, 50000
NumAggregators	integer >= 1	<b>0 (one file per compute node)</b> , MPI_Size/2, ... , 2, (N-to-1) 1
AggregatorRatio	integer >= 1	not used unless set, MPI_Size/N must be an integer value
OpenTimeoutSecs	float	<b>0</b> , 10.0, 5
BeginStepPollingFrequency-Secs	float	<b>1</b> , 10.0
StatsLevel	integer, 0 or 1	<b>1</b> , 0
StatsBlockSize	integer > 0	<b>a very big number</b> , 1073741824 for blocks with 1M elements
NodeLocal	string On/Off	<b>Off</b> , On
Node-Local	string On/Off	<b>Off</b> , On
BurstBufferPath	string	<b>“”</b> , /mnt/bb/norbert, /ssd
BurstBufferDrain	string On/Off	<b>On</b> , Off
BurstBufferVerbose	integer, 0-2	<b>0</b> , 1, 2
StreamReader	string On/Off	On, <b>Off</b>

Only file transport types are supported. Optional parameters for `IO::AddTransport` or in runtime config file transport field:

**Transport type: File**

Key	Value Format	Default and Examples
Library	string	<b>POSIX</b> (UNIX), <b>FStream</b> (Windows), stdio, IME

The IME transport directly reads and writes files stored on DDN’s IME burst buffer using the IME native API. To use the IME transport, IME must be available on the target system and ADIOS2 needs to be configured with `ADIOS2_USE_IME`. By default, data written to the IME is automatically flushed to the parallel filesystem at every `EndStep()` call. You can disable this automatic flush by setting the transport parameter `SyncToPFS` to `OFF`.

## 10.3 BP3

The BP3 Engine writes and reads files in ADIOS2 native binary-pack (bp) format. BP files are backwards compatible with ADIOS1.x and have the following structure given a “name” string passed as the first argument of `IO::Open`:

```
adios2::Engine bpFile = io.Open("name", adios2::Mode::Write);
```

will generate:



```
% collective metadata file
name.bp

% data directory and files
name.bp.dir/
    name.bp.0
    name.bp.1
    ...
    name.bp.M
```

**Note:** BP3 file names are compatible with the Unix (/) and Windows (\\) file system naming convention for directories and files.

**Caution:** The default BP3 engine will check if the .bp is the extension of the first argument of `IO::Open` and will add .bp and .bp.dir if not.

This engine allows the user to fine tune the buffering operations through the following optional parameters:

1. **Profile:** turns ON/OFF profiling information right after a run
2. **ProfileUnits:** set profile units according to the required measurement scale for intensive operations
3. **CollectiveMetadata:** turns ON/OFF forming collective metadata during run (used by large scale HPC applications)
4. **Threads:** number of threads provided from the application for buffering, use this for very large variables in data size
5. **InitialBufferSize:** initial memory provided for buffering (minimum is 16Kb)
6. **BufferGrowthFactor:** exponential growth factor for initial buffer > 1, default = 1.05.
7. **MaxBufferSize:** maximum allowable buffer size (must be larger than 16Kb). If too large adios2 will throw an exception.
8. **FlushStepsCount:** users can select how often to produce the more expensive collective metadata file in terms of steps: default is 1. Increase to reduce adios2 collective operations footprint, with the trade-off of reducing checkpoint frequency. Buffer size will increase until first steps count if **MaxBufferSize** is not set.
9. **NumAggregators** (or **SubStreams**): Users can select how many sub-files (M) are produced during a run, ranges between 1 and the number of mpi processes from **MPI\_Size** (N), adios2 will internally aggregate data buffers (N-to-M) to output the required number of sub-files. Default is 0, which will let adios2 to group processes per shared-memory-access (i.e. one per compute node) and use one process per node as an aggregator. If **NumAggregators** is larger than the number of processes then it will be set to the number of processes.
10. **AggregatorRatio:** An alternative option to **NumAggregators** to pick every Nth process as aggregator. An integer divider of the number of processes is required, otherwise a runtime exception is thrown.
11. **Node-Local:** For distributed file system. Every writer process must make sure the .bp/ directory is created on the local file system. Required for using local disk/SSD/NVMe in a cluster.

Key	Value Format	Default and Examples
Profile	string On/Off	<b>On</b> , Off
ProfileUnits	string	<b>Microseconds</b> , Milliseconds, Seconds, Minutes, Hours
CollectiveMetadata	string On/Off	<b>On</b> , Off
Threads	integer > 1	<b>1</b> , 2, 3, 4, 16, 32, 64
InitialBufferSize	float+units >= 16Kb	<b>16Kb</b> , 10Mb, 0.5Gb
MaxBufferSize	float+units >= 16Kb	<b>at EndStep</b> , 10Mb, 0.5Gb
BufferGrowthFactor	float > 1	<b>1.05</b> , 1.01, 1.5, 2
FlushStepsCount	integer > 1	<b>1</b> , 5, 1000, 50000
NumAggregators	integer >= 1	<b>0 (one file per compute node)</b> , <code>MPI_Size/2, ... , 2, (N-to-1)</code> 1
AggregatorRatio	integer >= 1	not used unless set, <code>MPI_Size/N</code> must be an integer value
Node-Local	string On/Off	<b>Off</b> , On

Only file transport types are supported. Optional parameters for `IO::AddTransport` or in runtime config file transport field:

**Transport type: File**

Key	Value Format	Default and Examples
Library	string	<b>POSIX</b> (UNIX), <b>FStream</b> (Windows), stdio, IME

## 10.4 HDF5

In ADIOS2, the default engine for reading and writing HDF5 files is called “*HDF5*”. To use this engine, you can either specify it in your xml config file, with tag `<engine type=HDF5>` or, set it in client code. For example, here is how to create a hdf5 reader:

```
adios2::IO h5IO = adios.DeclareIO("SomeName");
h5IO.SetEngine("HDF5");
adios2::Engine h5Reader = h5IO.Open(filename, adios2::Mode::Read);
```

To read back the h5 files generated with VDS to ADIOS2, one can use the HDF5 engine. Please make sure you are using the HDF5 library that has version greater than or equal to 1.11 in ADIOS2.

The h5 file generated by ADIOS2 has two levels of groups: The top Group, / and its subgroups: `Step0 ... StepN`, where N is number of steps. All datasets belong to the subgroups.

Any other h5 file can be read back to ADIOS as well. To be consistent, when reading back to ADIOS2, we assume a default Step0, and all datasets from the original h5 file belong to that subgroup. The full path of a dataset (from the original h5 file) is used when represented in ADIOS2.

We can pass options to HDF5 API from ADIOS xml configuration. Currently we support `CollectionIO` (default false), and chunk specifications. The chunk specification uses space to separate values, and by default, if a valid `H5ChunkDim` exists, it applies to all variables, unless `H5ChunkVar` is specified. Examples:

```
<parameter key="H5CollectiveMPIIO" value="yes"/>
<parameter key="H5ChunkDim" value="200 200"/>
<parameter key="H5ChunkVar" value="VarName1 VarName2"/>
```

We suggest to read HDF5 documentation before applying these options.

After the subfile feature is introduced in HDF5 version 1.14, the ADIOS2 HDF5 engine will use subfiles as the default h5 format as it improves I/O in general (for example, see <https://escholarship.org/uc/item/6fs7s3jb>)

To use the subfile feature, client needs to support MPI\_Init\_thread with MPI\_THREAD\_MULTIPLE.

Useful parameters from the HDF library to tune subfiles are:

```
H5FD_SUBFILING_IOC_PER_NODE (num of subfiles per node)
    set H5FD_SUBFILING_IOC_PER_NODE to 0 if the regular h5 file is preferred, before using
    ↪ ADIOS2 HDF5 engine.
H5FD_SUBFILING_STRIPE_SIZE
H5FD_IOC_THREAD_POOL_SIZE
```

## 10.5 SST Sustainable Staging Transport

In ADIOS2, the Sustainable Staging Transport (SST) is an engine that allows direct connection of data producers and consumers via the ADIOS2 write/read APIs. This is a classic streaming data architecture where the data passed to ADIOS on the write side (via Put() deferred and sync, and similar calls) is made directly available to a reader (via Get(), deferred and sync, and similar calls).

SST is designed for use in HPC environments and can take advantage of RDMA network interconnects to speed the transfer of data between communicating HPC applications; however, it is also capable of operating in a Wide Area Networking environment over standard sockets. SST supports full MxN data distribution, where the number of reader ranks can differ from the number of writer ranks. SST also allows multiple reader cohorts to get access to a writer's data simultaneously.

To use this engine, you can either specify it in your xml config file, with tag `<engine type=SST>` or, set it in client code. For example, here is how to create an SST reader:

```
adios2::IO sstIO = adios.DeclareIO("SomeName");
sstIO.SetEngine("SST");
adios2::Engine sstReader = sstIO.Open(filename, adios2::Mode::Read);
```

and a sample code for SST writer is:

```
adios2::IO sstIO = adios.DeclareIO("SomeName");
sstIO.SetEngine("SST");
adios2::Engine sstWriter = sstIO.Open(filename, adios2::Mode::Write);
```

The general goal of ADIOS2 is to ease the conversion of a file-based application to instead use a non-file streaming interconnect, for example, data producers such as computational physics codes and consumers such as analysis applications. However, there are some uses of ADIOS2 APIs that work perfectly well with the ADIOS2 file engines, but which will not work or will perform badly with streaming. For example, SST is based upon the “*step*” concept and ADIOS2 applications that use SST must call BeginStep() and EndStep(). On the writer side, the Put() calls between BeginStep and EndStep are the unit of communication and represent the data that will be available between the corresponding Begin/EndStep calls on the reader.

Also, it is recommended that SST-based applications not use the ADIOS2 Get() sync method unless there is only one data item to be read per step. This is because SST implements MxN data transfer (and avoids having to deliver all data to every reader), by queueing data on the writer ranks until it is known which reader rank requires it. Normally this data fetch stage is initiated by PerformGets() or EndStep(), both of which fulfill any pending Get() deferred operations. However, unlike Get() deferred, the semantics of Get() sync require the requested data to be fetched from the writers before the call can return. If there are multiple calls to Get() sync per step, each one may require a communication with many writers, something that would have only had to happen once if Get() deferred were used instead. Thus the use of Get() sync is likely to incur a substantial performance penalty.

On the writer side, depending upon the chosen data marshaling option there may be some (relatively small) performance differences between `Put()` sync and `Put()` deferred, but they are unlikely to be as substantial as between `Get()` sync and `Get()` deferred.

Note that SST readers and writers do not necessarily move in lockstep, but depending upon the queue length parameters and queueing policies specified, differing reader and writer speeds may cause one or the other side to wait for data to be produced or consumed, or data may be dropped if allowed by the queueing policy. However, steps themselves are atomic and no step will be partially dropped, delivered to a subset of ranks, or otherwise divided.

The SST engine allows the user to customize the streaming operations through the following optional parameters:

1. `RendezvousReaderCount`: Default **1**. This integer value specifies the number of readers for which the writer should wait before the writer-side `Open()` returns. The default of 1 implements an ADIOS1/flexpath style “rendezvous”, in which an early-starting reader will wait for the writer to start, or vice versa. A number >1 will cause the writer to wait for more readers and a value of 0 will allow the writer to proceed without any readers present. This value is interpreted by SST Writer engines only.

2. `RegistrationMethod`: Default **“File”**. By default, SST reader and writer engines communicate network contact information via files in a shared filesystem. Specifically, the “filename” parameter in the `Open()` call is interpreted as a path which the writer uses as the name of a file to which contact information is written, and from which a reader will attempt to read contact information. As with other file-based engines, file creation and access is subject to the usual considerations (directory components are interpreted, but must exist and be traversable, writer must be able to create the file and the reader must be able to read it). Generally the file so created will exist only for as long as the writer keeps the stream `Open()`, but abnormal process termination may leave “stale” files in those locations. These stray “.sst” files should be deleted to avoid confusing future readers. SST also offers a **“Screen”** registration method in which writers and readers send their contact information to, and read it from, `stdout` and `stdin` respectively. The “screen” registration method doesn’t support batch mode operations in any way, but may be useful when manually starting jobs on machines in a WAN environment that don’t share a filesystem. A future release of SST will also support a **“Cloud”** registration method where contact information is registered to and retrieved from a network-based third-party server so that both the shared filesystem and interactivity can be avoided. This value is interpreted by both SST Writer and Reader engines.

3. `QueueLimit`: Default **0**. This integer value specifies the number of steps which the writer will allow to be queued before taking specific action (such as discarding data or waiting for readers to consume the data). The default value of 0 is interpreted as no limit. This value is interpreted by SST Writer engines only.

4. `QueueFullPolicy`: Default **“Block”**. This value controls what policy is invoked if a non-zero `QueueLimit` has been specified and new data would cause the queue limit to be reached. Essentially, the **“Block”** option ensures data will not be discarded and if the queue fills up the writer will block on `EndStep` until the data has been read. If there is one active reader, `EndStep` will block until data has been consumed off the front of the queue to make room for newly arriving data. If there is more than one active reader, it is only removed from the queue when it has been read by all readers, so the slowest reader will dictate progress. **NOTE THAT THE NO READERS SITUATION IS A SPECIAL CASE**: If there are no active readers, new timesteps are considered to have completed their active queueing immediately upon submission. They may be retained in the “reserve queue” if the `ReserveQueueLimit` is non-zero. However, if that `ReserveQueueLimit` parameter is zero, timesteps submitted when there are no active readers will be immediately discarded.

Besides **“Block”**, the other acceptable value for `QueueFullPolicy` is **“Discard”**. When **“Discard”** is specified, and an `EndStep` operation would add more than the allowed number of steps to the queue, some step is discarded. If there are no current readers connected to the stream, the *oldest* data in the queue is discarded. If there are current readers, then the *newest* data (I.E. the just-created step) is discarded. (The differential treatment is because SST sends metadata for each step to the readers as soon as the step is accepted and cannot reliably prevent that use of that data without a costly all-to-all synchronization operation. Discarding the *newest* data instead is less satisfying, but has a similar long-term effect upon the set of steps delivered to the readers.) This value is interpreted by SST Writer engines only.

5. `ReserveQueueLimit`: Default **0**. This integer value specifies the number of steps which the writer will keep in the queue for the benefit of late-arriving readers. This may consist of timesteps that have already been consumed by any readers, as well as timesteps that have not yet been consumed. In some sense this is target queue minimum size, while `QueueLimit` is a maximum size. This value is interpreted by SST Writer engines only.

6. **DataTransport**: Default **varies**. This string value specifies the underlying network communication mechanism to use for exchanging data in SST. Generally this is chosen by SST based upon what is available on the current platform. However, specifying this engine parameter allows overriding SST's choice. Current allowed values are **"UCX"**, **"MPI"**, **"RDMA"**, and **"WAN"**. (**ib** and **fabric** are accepted as equivalent to **RDMA** and **evpath** is equivalent to **WAN**.) Generally both the reader and writer should be using the same network transport, and the network transport chosen may be dictated by the situation. For example, the RDMA transport generally operates only between applications running on the same high-performance interconnect (e.g. on the same HPC machine). If communication is desired between applications running on different interconnects, the Wide Area Network (WAN) option should be chosen. This value is interpreted by both SST Writer and Reader engines.

7. **WANDataTransport**: Default **sockets**. If the SST **DataTransport** parameter is **"WAN"**, this string value specifies the EVPath-level data transport to use for exchanging data. The value must be a data transport known to EVPath, such as **"sockets"**, **"enet"**, or **"ib"**. Generally both the reader and writer should be using the same EVPath-level data transport. This value is interpreted by both SST Writer and Reader engines.

8. **ControlTransport**: Default **tcp**. This string value specifies the underlying network communication mechanism to use for performing control operations in SST. SST can be configured to standard TCP sockets, which are very reliable and efficient, but which are limited in their scalability. Alternatively, SST can use a reliable UDP protocol, that is more scalable, but as of ADIOS2 Release 2.4.0 still suffers from some reliability problems. (**sockets** is accepted as equivalent to **tcp** and **udp**, **rudp**, and **enet** are equivalent to **scalable**. Generally both the reader and writer should be using the same control transport. This value is interpreted by both SST Writer and Reader engines.

9. **NetworkInterface**: Default **NULL**. In situations in which there are multiple possible network interfaces available to SST, this string value specifies which should be used to generate SST's contact information for writers. Generally this should *NOT* be specified except for narrow sets of circumstances. It has no effect if specified on Reader engines. If specified, the string value should correspond to a name of a network interface, such as are listed by commands like **"netstat -i"**. For example, on most Unix systems, setting the NetworkInterface parameter to **"lo"** (or possibly **"lo0"**) will result in SST generating contact information that uses the network address associated with the loopback interface (127.0.0.1). This value is interpreted by only by the SST Writer engine.

10. **ControlInterface**: Default **NULL**. This value is similar to the NetworkInterface parameter, but only applies to the SST layer which does messaging for control (open, close, flow and timestep management, but not actual data transfer). Generally the NetworkInterface parameter can be used to control this, but that also applies to the Data Plane. Use ControlInterface in the event of conflicting specifications.

11. **DataInterface**: Default **NULL**. This value is similar to the NetworkInterface parameter, but only applies to the SST layer which does messaging for data transfer, not control (open, close, flow and timestep management). Generally the NetworkInterface parameter can be used to control this, but that also applies to the Control Plane. Use DataInterface in the event of conflicting specifications. In the case of the RDMA data plane, this parameter controls the libfabric interface choice.

12. **FirstTimestepPrecious**: Default **FALSE**. FirstTimestepPrecious is a boolean parameter that affects the queueing of the first timestep presented to the SST Writer engine. If FirstTimestepPrecious is **TRUE**, then the first timestep is effectively never removed from the output queue and will be presented as a first timestep to any reader that joins at a later time. This can be used to convey run parameters or other information that every reader may need despite joining later in a data stream. Note that this queued first timestep does count against the QueueLimit parameter above, so if a QueueLimit is specified, it should be a value larger than 1. Further note while specifying this parameter guarantees that the preserved first timestep will be made available to new readers, other reader-side operations (like requesting the LatestAvailable timestep in Engine parameters) might still cause the timestep to be skipped. This value is interpreted by only by the SST Writer engine.

13. **AlwaysProvideLatestTimestep**: Default **FALSE**. AlwaysProvideLatestTimestep is a boolean parameter that affects what of the available timesteps will be provided to the reader engine. If AlwaysProvideLatestTimestep is **TRUE**, then if there are multiple timesteps available to the reader, older timesteps will be skipped and the reader will see only the newest available upon BeginStep. This value is interpreted by only by the SST Reader engine.

14. **OpenTimeoutSecs**: Default **60**. OpenTimeoutSecs is an integer parameter that specifies the number of seconds SST is to wait for a peer connection on Open(). Currently this is only implemented on the Reader side of SST, and is

a timeout for locating the contact information file created by Writer-side Open, not for completing the entire Open() handshake. Currently value is interpreted by only by the SST Reader engine.

15. **SpeculativePreloadMode**: Default **AUTO**. In some circumstances, SST eagerly sends all data from writers to every readers without first waiting for read requests. Generally this improves performance if every reader needs all the data, but can be very detrimental otherwise. The value **AUTO** for this engine parameter instructs SST to apply its own heuristic for determining if data should be eagerly sent. The value **OFF** disables this feature and the value **ON** causes eager sending regardless of heuristic. Currently SST's heuristic is simple. If the size of the reader cohort is less than or equal to the value of the **SpecAutoNodeThreshold** engine parameter (Default value 1), eager sending is initiated. Currently value is interpreted by only by the SST Reader engine.

16. **SpecAutoNodeThreshold**: Default **1**. If the size of the reader cohort is less than or equal to this value *and* the **SpeculativePreloadMode** parameter is **AUTO**, SST will initiate eager data sending of all data from each writer to all readers. Currently value is interpreted by only by the SST Reader engine.

17. **StepDistributionMode**: Default **"AllToAll"**. This value controls how steps are distributed, particularly when there are multiple readers. By default, the value is **"AllToAll"**, which means that all timesteps are to be delivered to all readers (subject to discard rules, etc.). In other distribution modes, this is not the case. For example, in **"RoundRobin"**, each step is delivered only to a single reader, determined in a round-robin fashion based upon the number of readers who have opened the stream at the time the step is submitted. In **"OnDemand"** each step is delivered to a single reader, but only upon request (with a request being initiated by the reader doing **BeginStep()**). Normal reader-side rules (like **BeginStep** timeouts) and writer-side rules (like queue limit behavior) apply.

Key		Value Format	Default and Examples
RendezvousReaderCount	RegistrationMethod	integer string	<b>1 File</b> , Screen <b>0</b> (no queue limits)
QueueLimit	QueueFullPolicy	integer string in-	<b>Block</b> , Discard <b>0</b> (no queue limits) <b>de-</b>
DataTransport	WANDataTransport	integer string string	<b>fault varies by platform</b> , UCX, MPI,
MarshalMethod	NetworkInterface	string string	RDMA, WAN <b>sockets</b> , enet, ib <b>TCP</b> ,
ControlInterface	DataInterface	string string	Scalable <b>BP5</b> , BP, FFS <b>NULL NULL</b>
AlwaysProvideLatestTimestep	OpenTimeoutSecs	string boolean	<b>NULL FALSE</b> , true, no, yes <b>FALSE</b> ,
SpeculativePreloadMode	SpecAutoNodeThreshold	boolean integer	true, no, yes <b>60 AUTO</b> , ON, OFF <b>1</b>
		string integer	

## 10.6 SSC Strong Staging Coupler

The SSC engine is designed specifically for strong code coupling. Currently SSC only supports fixed IO pattern, which means once the first step is finished, users are not allowed to write or read a data block with a *start* and *count* that have not been written or read in the first step. SSC uses a combination of one sided MPI and two sided MPI methods. In any cases, all user applications are required to be launched within a single mpirun or mpiexec command, using the MPMD mode.

The SSC engine takes the following parameters:

1. **OpenTimeoutSecs**: Default **10**. Timeout in seconds for opening a stream. The SSC engine's open function will block until the **RendezvousAppCount** is reached, or timeout, whichever comes first. If it reaches the timeout, SSC will throw an exception.
2. **Threading**: Default **False**. SSC will use threads to hide the time cost for metadata manipulation and data transfer when this parameter is set to **true**. SSC will check if MPI is initialized with multi-thread enabled, and if not, then SSC will force this parameter to be **false**. Please do NOT enable threading when multiple I/O streams are opened in an application, as it will cause unpredictable errors. This parameter is only effective when writer definitions and reader selections are NOT locked. For cases definitions and reader selections are locked, SSC has a more optimized way to do data transfers, and thus it will not use this parameter.

Key	Value Format	Default and Examples
OpenTimeoutSecs	integer	<b>10</b> , 2, 20, 200
Threading	bool	<b>false</b> , true

## 10.7 DataMan for Wide Area Network Data Staging

The DataMan engine is designed for data staging over the wide area network. It is supposed to be used in cases where a few writers send data to a few readers over long distance.

DataMan supports compression operators such as ZFP lossy compression and BZip2 lossless compression. Please refer to the operator section for usage.

The DataMan engine takes the following parameters:

1. **IPAddress**: No default value. The IP address of the host where the writer application runs. This parameter is compulsory in wide area network data staging.
2. **Port**: Default **50001**. The port number on the writer host that will be used for data transfers.
3. **Timeout**: Default **5**. Timeout in seconds to wait for every send / receive operation. Packages not sent or received within this time are considered lost.
4. **RendezvousReaderCount**: Default **1**. This integer value specifies the number of readers for which the writer should wait before the writer-side Open() returns. By default, an early-starting writer will wait for the reader to start, or vice versa. A number >1 will cause the writer to wait for more readers, and a value of 0 will allow the writer to proceed without any readers present. This value is interpreted by DataMan Writer engines only.
5. **Threading**: Default **true** for reader, **false** for writer. Whether to use threads for send and receive operations. Enabling threading will cause extra overhead for managing threads and buffer queues, but will improve the continuity of data steps for readers, and help overlap data transfers with computations for writers.
6. **TransportMode**: Default **fast**. Only DataMan writers take this parameter. Readers are automatically synchronized at runtime to match writers' transport mode. The fast mode is optimized for latency-critical applications. It enforces readers to only receive the latest step. Therefore, in cases where writers are faster than readers, readers will skip some data steps. The reliable mode ensures that all steps are received by readers, by sacrificing performance compared to the fast mode.
7. **MaxStepBufferSize**: Default **128000000**. In order to bring down the latency in wide area network staging use cases, DataMan uses a fixed receiver buffer size. This saves an extra communication operation to sync the buffer size for each step, before sending actual data. The default buffer size is 128 MB, which is sufficient for most use cases. However, in case 128 MB is not enough, this parameter must be set correctly, otherwise DataMan will fail.

Key	Value Format	Default and Examples
IPAddress	string	N/A, 22.195.18.29
Port	integer	<b>50001</b> , 22000, 33000
Timeout	integer	<b>5</b> , 10, 30
RendezvousReaderCount	integer	<b>1</b> , 0, 3
Threading	bool	<b>true</b> for reader, <b>false</b> for writer
TransportMode	string	<b>fast</b> , reliable
MaxStepBufferSize	integer	<b>128000000</b> , 512000000, 1024000000

## 10.8 DataSpaces

The DataSpaces engine for ADIOS2 is experimental. DataSpaces is an asynchronous I/O transfer method within ADIOS that enables low-overhead, high-throughput data extraction from a running simulation. DataSpaces is designed for use in HPC environments and can take advantage of RDMA network interconnects to speed the transfer of data between communicating HPC applications. DataSpaces supports full MxN data distribution, where the number of reader ranks can differ from the number of writer ranks. In addition, this engine supports multiple reader and writer applications, which must be distinguished by unique values of AppID for different applications. It can be set in the xml config file with tag `<parameter key="AppID" value="2"/>`. The value should be unique for each applications or clients.

To use this engine, you can either specify it in your xml config file, with tag `<engine type=DATASPACES>` or, set it in client code. For example, here is how to create an DataSpaces reader:

```
adios2::IO dspacesIO = adios.DeclareIO("SomeName");
dspacesIO.SetEngine("DATASPACES");
adios2::Engine dspacesReader = dspacesIO.Open(filename, adios2::Mode::Read);
```

and a sample code for DataSpaces writer is:

```
adios2::IO dspacesIO = adios.DeclareIO("SomeName");
dspacesIO.SetEngine("DATASPACES");
adios2::Engine dspacesWriter = dspacesIO.Open(filename, adios2::Mode::Write);
```

To make use of the DataSpaces engine, an application job needs to also run the `dataspaces_server` component together with the application. The server should be configured and started before the application as a separate job in the system. For example:

```
aprun -n $SPROC ./dataspaces_server -s $SPROC &> log.server &
```

The variable `$SPROC` represents the number of server instances to run. The `&` character at the end of the line would place the `aprun` command in the background, and will allow the job script to continue and run the other applications. The server processes produce a configuration file, i.e., `conf.0` that is used by the application to connect to the servers. This file contains identifying information of the master server, which coordinates the client registration and discovery process. The job script should wait for the servers to start-up and produce the `conf.0` file before starting the client application processes.

The server also needs a user configuration read from a text file called `dataspaces.conf`. How many output timesteps of the same dataset (called versions) should be kept in the server's memory and served to readers should be specified in the file. If this file does not exist in the current directory, the server will assume default values (only 1 timestep stored). .. code-block:

```
## Config file for DataSpaces
max_versions = 5
lock_type = 3
```

The `dataspaces_server` module is a stand-alone service that runs independently of a simulation on a set of dedicated nodes in the staging area. It transfers data from the application through RDMA, and can save it to local storage system, e.g., the Lustre file system, stream it to remote sites, e.g., auxilliary clusters, or serve it directly from the staging area to other applications. One instance of the `dataspaces_server` can service multiple applications in parallel. Further, the server can run in cooperative mode (i.e., multiple instances of the server cooperate to service the application in parallel and to balance load). The `dataspaces_server` receives notification messages from the transport method, schedules the requests, and initiates the data transfers in parallel. The server schedules and prioritizes the data transfers while the simulation is computing in order to overlap data transfers with computations, to maximize data throughput, and to minimize the overhead on the application.



## 10.9 Inline for zero-copy

The Inline engine provides in-process communication between the writer and reader, avoiding the copy of data buffers.

This engine is focused on the  $N \rightarrow N$  case:  $N$  writers share a process with  $N$  readers, and the analysis happens ‘inline’ without writing the data to a file or copying to another buffer. It is similar to the streaming SST engine, since analysis must happen per step.

To use this engine, you can either add `<engine type=Inline>` to your XML config file, or set it in your application code:

```
adios2::IO io = adios.DeclareIO("ioName");
io.SetEngine("Inline");
adios2::Engine inlineWriter = io.Open("inline_write", adios2::Mode::Write);
adios2::Engine inlineReader = io.Open("inline_read", adios2::Mode::Read);
```

Notice that unlike other engines, the reader and writer share an IO instance. Both the writer and reader must be opened before either tries to call `BeginStep()/PerformPuts()/PerformGets()`. There must be exactly one writer, and exactly one reader.

For successful operation, the writer will perform a step, then the reader will perform a step in the same process. When the reader starts its step, the only data it has available is that written by the writer in its process. The reader then can retrieve whatever data was written by the writer by using the double-pointer `Get` call:

```
void Engine::Get<T>(Variable<T>, T**) const;
```

This version of `Get` is only used for the inline engine. See the example below for details.

**Note:** Since the inline engine does not copy any data, the writer should avoid changing the data before the reader has read it.

Typical access pattern:

```
// ... Application data generation

inlineWriter.BeginStep();
inlineWriter.Put(var, in_data); // always use deferred mode
inlineWriter.EndStep();
// Unlike other engines, data should not be reused at this point (since ADIOS
// does not copy the data), though ADIOS cannot enforce this.
// Must wait until reader is finished using the data.

inlineReader.BeginStep();
double* out_data;
inlineReader.Get(var, &data);
// Now in_data == out_data.
inlineReader.EndStep();
```

## 10.10 Null

The Null Engine performs no internal work and no I/O. It was created for testing applications that have ADIOS2 output in it by turning off the I/O easily. The runtime difference between a run with the Null engine and another engine tells us the IO overhead of that particular output with that particular engine.

```
adios2::IO io = adios.DeclareIO("Output");  
io.SetEngine("Null");
```

or using the XML config file:

```
<adios-config>  
  <io name="Output">  
    <engine type="Null">  
    </engine>  
  </io>  
</adios-config>
```

Although there is a reading engine as well, which will not fail, any variable/attribute inquiry returns *nullptr* and any subsequent Get() calls will throw an exception in C++/Python or return an error in C/Fortran.

Note that there is also a *Null transport* that can be used by a BP engine instead of the default *File transport*. In that case, the BP engine will perform all internal work including buffering and aggregation but no data will be output at all. A run like this can be used to assess the overhead of the internal work of the BP engine.

```
adios2::IO io = adios.DeclareIO("Output");  
io.SetEngine("BP5");  
io.AddTransport("Null", {});
```

or using the XML config file

```
<adios-config>  
  <io name="Output">  
    <engine type="BP5">  
    </engine>  
    <transport type="Null">  
    </transport>  
  </io>  
</adios-config>
```

## 10.11 Plugin Engine

For details on using the Plugin Engine, see the *Plugins* documentation.

## SUPPORTED OPERATORS

The Operator abstraction allows ADIOS2 to act upon the user application data, either from a `adios2::Variable` or a set of Variables in an `adios2::IO` object. Current supported operations are:

1. Data compression/decompression, lossy and lossless.

This section provides a description of the supported operators in ADIOS2 and their specific parameters to allow extra-control from the user. Parameters are passed in key-value pairs for:

1. Operator general supported parameters.
2. Operator specific supported parameters.

Parameters are passed at:

1. Compile time: using the second parameter of the method `ADIOS2::DefineOperator`
2. *Runtime Configuration Files* in the *ADIOS* component.

### 11.1 CompressorZFP

The `CompressorZFP` Operator is compressor that uses a lossy but optionally error-bounded compression to achieve high compression ratios.

ZFP provides compressed-array classes that support high throughput read and write random access to individual array elements. ZFP also supports serial and parallel (OpenMP and CUDA) compression of whole arrays, e.g., for applications that read and write large data sets to and from disk.

ADIOS2 provides a `CompressorZFP` operator that lets you compress an decompress variables. Below there is an example of how to invoke `CompressorZFP` operator:

```
adios2::IO io = adios.DeclareIO("Output");
auto ZFP0p    = adios.DefineOperator("CompressorZFP", adios2::ops::LossyZFP);

auto var_r32 = io.DefineVariable<float>("r32", shape, start, count);
var_r32.AddOperation(ZFP0p, {{adios2::ops::zfp::key::rate, rate}});
```

### 11.1.1 CompressorZFP Specific parameters

The CompressorZFP operator accepts the following operator specific parameters:

CompressorZFP available parameters	
accuracy	Fixed absolute error tolerance
rate	Fixed number of bits in a compression unit
precision	Fixed number of uncompressed bits per value
backend	Backend device: cuda omp serial

### 11.1.2 CompressorZFP Execution Policy

CompressorZFP can run in multiple backend devices: GPUs (CUDA), OpenMP, and in the host CPU. By default CompressorZFP will choose its backend following the above order upon the availability of the device adapter.

Exceptionally, if its corresponding ADIOS2 variable contains a CUDA memory address, this is a CUDA buffer, the CUDA backend will be called if available.

In any case, the user can manually set the backend using the ZFPOperator specific parameter `backend`.

## 11.2 Plugin Operator

For details on using the Plugin Operator, see the [Plugins](#) documentation.

## 11.3 Encryption

The Encryption Operator uses the [Plugins](#) interface. This operator uses [libsodium](#) for encrypting and decrypting data. If ADIOS can find libsodium at configure time, this plugin will be built.

This operator will generate a secret key and encrypts the data with the key and a nonce as described in the [libsodium secret key cryptography docs](#). The key is saved to the specified `SecretKeyFile` and will be used for decryption. The key should be kept confidential since it is used to both encrypt and decrypt the data.

Parameters to use with the Encryption operator:

Key	Value Format	Explanation
PluginName	string	Required. Name to refer to plugin, e.g., <code>MyOperator</code>
PluginLibrary	string	Required. Name of shared library, <code>EncryptionOperator</code>
SecretKeyFile	string	Required. Path to secret key file

## FULL APIS

---

**Note:** Application developers who desire fine-grained control of IO tasks should use the Full APIs. In simple cases (e.g. reading a file for analysis, interactive Python, or saving some data for a small project or tests) please refer to the *High-Level APIs*.

---

Currently ADIOS2 support bindings for the following languages and their minimum standards:

Language	Standard	Interface
C++	11/newer	<code>#include adios2.h</code>
	older	use C bindings
C	99	<code>#include adios2_c.h</code>
Fortran	90	use <code>adios2</code>
Python	2.7	<code>import adios2</code>
	3	<code>import adios2</code>

---

**Tip:** Prefer the C++11 bindings if your application C++ compiler supports the 2011 (or later) standard. For code using previous C++ standards (98 or 03) use the C bindings for ABI compatibility.

---

**Caution:** Statically linked libraries (\*.a) might result in conflicting ABIs between an older C++ project, the C bindings, and the adios native C++11 library. Test to make sure it works for your platform.

The current interaction flow for each language binding API with the ADIOS2 library is specified as follows

The following sections provide a summary of the API calls on each language and links to Write and Read examples to put it all together.

### 12.1 C++11 bindings

**Caution:** DO NOT use the clause `using namespace adios2` in your code. This is in general a bad practices that creates potential name conflicts. Always use `adios2::` explicitly, e.g. `adios2::ADIOS`, `adios2::IO`.

**Tip:** Prefer the C++11 bindings to take advantage of added functionality (*e.g.* move semantics, lambdas, etc.). If you must use an older C++ standard (98 or 03) to avoid application binary interface (ABI) incompatibilities use the C bindings.

---

### 12.1.1 ADIOS2 components classes

ADIOS2 C++ bindings objects are mapped 1-to-1 to the ADIOS components described in the [Components Overview](#) section. Only the `adios2::ADIOS` object is “owned” by the developer’s program using `adios2`, all other components are light-weight objects that point internally to a component that lives inside the `adios2::ADIOS` “factory” object.

```
C++11
adios2::ADIOS
adios2::IO
adios2::Variable<T>
adios2::Attribute<T>
adios2::Engine
adios2::Operator
```

The following section provides a summary of the available functionality for each class.

### 12.1.2 ADIOS class

class **ADIOS**

#### Public Functions

**ADIOS**(MPI\_Comm comm)

Starting point for MPI apps. Creates an *ADIOS* object. MPI Collective Operation as it call MPI\_Comm\_dup

##### Parameters

**comm** – defines domain scope from application

##### Throws

`std::invalid_argument` – if user input is incorrect

**ADIOS**(const std::string &configFile, MPI\_Comm comm)

Starting point for MPI apps. Creates an *ADIOS* object allowing a runtime config file. MPI collective and it calls MPI\_Comm\_dup and MPI\_Bcast to pass the configFile contents

##### Parameters

- **configFile** – runtime config file
- **comm** – defines domain scope from application

##### Throws

`std::invalid_argument` – if user input is incorrect

**ADIOS**(const std::string &configFile, MPI\_Comm comm, const std::string &hostLanguage)

extra constructor for R and other languages that use the public C++ API but has data in column-major. Pass “” for configfile if there is no config file. Last bool argument exist only to ensure matching this signature by

having different number of arguments. Supported languages are “R”, “Matlab”, “Fortran”, all these names mean the same thing: treat all arrays column-major e.g. `adios2::ADIOS(“”, comm, “Fortran”, false);`

**ADIOS**(const std::string &configFile)

Starting point for non-MPI serial apps. Creates an *ADIOS* object allowing a runtime config file.

**Parameters**

**configFile** – runtime config file

**Throws**

std::invalid\_argument – if user input is incorrect

**ADIOS**()

Starting point for non-MPI apps. Creates an *ADIOS* object

**Throws**

std::invalid\_argument – if user input is incorrect

**ADIOS**(const std::string &configFile, const std::string &hostLanguage)

extra constructor for R and other languages that use the public C++ API but has data in column-major. Pass “” for configfile if there is no config file. Last bool argument exist only to ensure matching this signature by having different number of arguments. Supported languages are “R”, “Matlab”, “Fortran”, all these names mean the same thing: treat all arrays column-major e.g. `adios2::ADIOS(“”, “Fortran”, false);`

explicit **operator bool**() const noexcept

object inspection true: valid object, false: invalid object

**ADIOS**(const *ADIOS*&) = delete

DELETED Copy Constructor. *ADIOS* is the only object that manages its own memory. Create a separate object for independent tasks

**ADIOS**(*ADIOS*&&) = default

default move constructor exists to allow for auto ad = ADIOS(...) initialization

**~ADIOS**() = default

MPI Collective calls MPI\_Comm\_free Uses RAII for all other members

*ADIOS* &**operator**=(const *ADIOS*&) = delete

copy assignment is forbidden for the same reason as copy constructor

*ADIOS* &**operator**=(*ADIOS*&&) = default

move assignment is allowed, though, to be consistent with move constructor

**IO DeclareIO**(const std::string name, const ArrayOrdering ArrayOrder = ArrayOrdering::Auto)

Declares a new *IO* class object

**Parameters**

**name** – unique *IO* name identifier within current *ADIOS* object

**Throws**

std::invalid\_argument – if *IO* with unique name is already declared

**Returns**

reference to newly created *IO* object inside current *ADIOS* object

**IO AtIO**(const std::string name)

Retrieve an existing *IO* object previously created with DeclareIO.

**Parameters**

**name** – *IO* unique identifier key in current *ADIOS* object

**Throws**

`std::invalid_argument` – if *IO* was not created with `DeclareIO`

**Returns**

if *IO* exists returns a reference to existing *IO* object inside *ADIOS*, else throws an exception.  
*IO* objects can't be invalid.

*Operator* **DefineOperator**(const std::string name, const std::string type, const Params &parameters = Params())

Defines an adios2 supported operator by its type.

**Parameters**

- **name** – unique operator name identifier within current *ADIOS* object
- **type** – supported ADIOS2 operator type: zfp, sz
- **parameters** – key/value parameters at the operator object level

**Throws**

`std::invalid_argument` – if adios2 can't support current operator due to missing dependency or unsupported type

**Returns**

*Operator* object

template<class **R**, class ...**Args**>

*Operator* **DefineOperator**(const std::string name, const std::function<*R*(*Args*...)> &function, const Params &parameters = Params())

Defines an adios2 supported operator by its type. Variadic template version for Operators of type Callback function with signatures supported in ADIOS2. For new signature support open an issue on github.

**Parameters**

- **name** – unique operator name within *ADIOS* object
- **function** – C++11 callable target
- **parameters** – key/value parameters at the operator level

**Throws**

`std::invalid_argument` – if adios2 can't support current operator due to missing dependency or unsupported type

**Returns**

*Operator* object

*Operator* **InquireOperator**(const std::string name)

Retrieve an existing *Operator* object in current *ADIOS* object

**Parameters**

**name** – *Operator* unique identifier key in current *ADIOS* object

**Returns**

object to an existing operator in current *ADIOS* object, *Operator* object is false if name is not found

void **FlushAll**()

Flushes all engines in write mode in all IOs created with the current *ADIOS* object. If no *IO* or *Engine* exist, it does nothing.

**Throws**

`std::runtime_error` – if any engine Flush fails



bool **RemoveIO**(const std::string name)

DANGER ZONE: removes a particular *IO*. This will effectively eliminate any parameter from the config.xml file

**Parameters**

**name** – io input name

**Returns**

true: *IO* was found and removed, false: *IO* not found and not removed

void **RemoveAllIOs**() noexcept

DANGER ZONE: removes all IOs created with DeclareIO. This will effectively eliminate any parameter from the config.xml file also.

void **EnterComputationBlock**() noexcept

Inform *ADIOS* about entering communication-free computation block in main thread. Useful when using Async *IO*

void **ExitComputationBlock**() noexcept

Inform *ADIOS* about exiting communication-free computation block in main thread. Useful when using Async *IO*

### 12.1.3 IO class

class **IO**

#### Public Functions

**IO**() = default

Empty (default) constructor, use it as a placeholder for future *IO* objects from ADIOS::IO functions. Can be used with STL containers.

**~IO**() = default

Use RAII

explicit **operator bool**() const noexcept

true: valid object, false: invalid object

std::string **Name**() const

Inspects *IO* name

**Returns**

name

bool **InConfigFile**() const

Checks if *IO* exists in a config file passed to *ADIOS* object that created this *IO*.

**Returns**

true: in config file, false: not in config file

void **SetEngine**(const std::string engineType)

Sets the engine type for current *IO* object.

**Parameters**

**engineType** – predefined engine type, default is bpfiler

void **SetParameter**(const std::string key, const std::string value)

Sets a single parameter. Overwrites value if key exists;.

**Parameters**

- **key** – parameter key
- **value** – parameter value

void **SetParameters**(const adios2::Params &parameters = adios2::Params())

Version that passes a map to fill out parameters initializer list = { “param1”, “value1” }, {“param2”, “value2”}, Replaces any existing parameter. Otherwise use SetParameter for adding new parameters.

**Parameters**

**parameters** – adios2::Params = std::map<std::string, std::string> key/value parameters

void **SetParameters**(const std::string &parameters)

Version that passes a single string to fill out many parameters. Replaces any existing parameter. initializer string = “param1=value1 , param2 = value2”.

void **ClearParameters**()

Remove all existing parameters. Replaces any existing parameter. initializer string = “param1=value1 , param2 = value2”.

adios2::Params **Parameters**() const

Return current parameters set from either SetParameters/SetParameter functions or from config XML for current *IO* object

**Returns**

string key/value map of current parameters (not modifiable)

size\_t **AddTransport**(const std::string type, const adios2::Params &parameters = adios2::Params())

Adds a transport and its parameters to current *IO*. Must be supported by current *EngineType*().

**Parameters**

- **type** – must be a supported transport type for a particular *Engine*. CAN’T use the keywords “Transport” or “transport”
- **parameters** – acceptable parameters for a particular transport

**Throws**

std::invalid\_argument – if type=transport

**Returns**

transportIndex handler

void **SetTransportParameter**(const size\_t transportIndex, const std::string key, const std::string value)

Sets a single parameter to an existing transport identified with a transportIndex handler from AddTransport. Overwrites existing parameter with the same key.

**Parameters**

- **transportIndex** – index handler from AddTransport
- **key** – parameter key
- **value** – parameter value

**Throws**

std::invalid\_argument – if transportIndex not valid, e.g. not a handler from AddTransport.

```
template<class T>
Variable<T> DefineVariable(const std::string &name, const Dims &shape = Dims(), const Dims &start =
                        Dims(), const Dims &count = Dims(), const bool constantDims = false)
```

Define a Variable<T> object within *IO*

#### Parameters

- **name** – unique variable identifier
- **shape** – global dimension
- **start** – local offset
- **count** – local dimension
- **constantDims** – true: shape, start, count won't change, false: shape, start, count will change after definition

#### Returns

Variable<T> object

```
template<class T>
Variable<T> InquireVariable(const std::string &name)
```

Retrieve a Variable object within current *IO* object

#### Parameters

**name** – unique variable identifier within *IO* object

#### Returns

if found Variable object is true and has functionality, else false and has no functionality

```
template<class T>
Attribute<T> DefineAttribute(const std::string &name, const T *data, const size_t size, const std::string
                        &variableName = "", const std::string separator = "/", const bool
                        allowModification = false)
```

Define attribute inside io. Array input version.

#### Parameters

- **name** – unique attribute identifier *IO* object or for a Variable if variableName is not empty (associated to a variable)
- **data** – pointer to user data
- **size** – number of data elements
- **variableName** – default is empty, if not empty attributes is associated to a variable
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.
- **allowModification** – true allows redefining existing attribute

#### Throws

std::invalid\_argument – if Attribute with unique name (in *IO* or Variable) is already defined

#### Returns

object reference to internal Attribute in *IO*

```
template<class T>
```

*Attribute*<T> **DefineAttribute**(const std::string &name, const T &value, const std::string &variableName = "", const std::string separator = "/", const bool allowModification = false)

Define single value attribute.

#### Parameters

- **name** – must be unique for the *IO* object or for a *Variable* if variableName is not empty (associated to a variable)
- **value** – single data value
- **variableName** – default is empty, if not empty attributes is associated to a variable
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.
- **allowModification** – true allows redefining existing attribute

#### Throws

std::invalid\_argument – if *Attribute* with unique name (in *IO* or *Variable*) is already defined

#### Returns

object reference to internal *Attribute* in *IO*

template<class T>

*Attribute*<T> **InquireAttribute**(const std::string &name, const std::string &variableName = "", const std::string separator = "/")

Retrieve an existing attribute.

#### Parameters

- **name** – must be unique for the *IO* object or for a *Variable* if variableName is not empty (associated to a variable)
- **variableName** – default is empty, if not empty attributes is expected to be associated to a variable
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.

#### Returns

object reference to internal *Attribute* in *IO*, object is false if *Attribute* is not found

bool **RemoveVariable**(const std::string &name)

DANGEROUS! Removes an existing *Variable* in current *IO* object. Might create dangling objects.

#### Parameters

**name** – unique *Variable* input

#### Returns

true: found and removed variable, false: not found, nothing to remove

void **RemoveAllVariables**()

DANGEROUS! Removes all existing variables in current *IO* object. Might create dangling objects.

bool **RemoveAttribute**(const std::string &name)

DANGEROUS! Removes an existing *Attribute* in current *IO* object. Might create dangling objects.

#### Parameters

**name** – unique *Attribute* identifier

**Returns**

true: found and removed attribute, false: not found, nothing to remove

void **RemoveAllAttributes**()

DANGEROUS! Removes all existing attributes in current *IO* object. Might create dangling objects.

*Engine* **Open**(const std::string &name, const Mode mode)

Open an *Engine* to start heavy-weight input/output operations.

**Parameters**

- **name** – unique engine identifier
- **mode** – adios2::Mode::Write, adios2::Mode::Read, adios2::Mode::ReadStreaming, or adios2::Mode::Append (BP4 only)

**Returns**

engine object

*Group* **InquireGroup**(char delimiter = '/')

Return a *Group* object for hierarchical reading.

**Parameters**

- **name** – starting path
- **a** – delimiter to separate groups in a string representation

**Returns**

*Group* object

*Engine* **Open**(const std::string &name, const Mode mode, MPI\_Comm comm)

Open an *Engine* to start heavy-weight input/output operations. This version allows passing a MPI communicator different from the one used in the *ADIOS* object constructor MPI Collective function as it calls MPI\_Comm\_dup

**Parameters**

- **name** – unique engine identifier within *IO*
- **mode** – adios2::Mode::Write, adios2::Mode::Read, or adios2::Mode::Append (BP4 only)
- **comm** – new communicator other than *ADIOS* object's communicator

**Returns**

engine object

void **FlushAll**()

Flushes all engines created with this *IO* with the Open function

std::map<std::string, Params> **AvailableVariables**(bool namesOnly = false)

Returns a map with variable information.

- key: variable name
- value: Params is a map<string,string>
  - key: “Type”, “Shape”, “AvailableStepsCount”, “Min”, “Max”, “SingleValue”

value: variable info value as string

**Parameters**

**namesOnly** – returns a map with the variable names but with no Parameters. Use this if you only need the list of variable names and call *VariableType()* and *InquireVariable()* on the names individually.

**Returns**

map<string, map<string, string>>

std::map<std::string, Params> **AvailableAttributes**(const std::string &variableName = "", const std::string separator = "/", const bool fullNameKeys = false)

Returns a map with available attributes information associated to a particular variableName

**Parameters**

- **variableName** – unique variable name associated with resulting attributes, if empty (default) return all attributes
- **separator** – optional name hierarchy separator (/, ::, \_, -, \, etc.)
- **fullNameKeys** – true: return full attribute names in keys, false (default): return attribute names relative to variableName

**Returns**

map:

std::string **VariableType**(const std::string &name) const

Inspects variable type. This function can be used in conjunction with MACROS in an else if (type == adios2::GetType<T>() ) { } loop

**Parameters**

**name** – unique variable name identifier in current *IO*

**Returns**

type as in adios2::GetType<T>() (e.g. “double”, “float”), empty std::string if variable not found

std::string **AttributeType**(const std::string &name) const

Inspects attribute type. This function can be used in conjunction with MACROS in an else if (type == adios2::GetType<T>() ) { } loop

**Parameters**

**name** – unique attribute name identifier in current *IO*

**Returns**

type as in adios2::GetType<T>() (e.g. “double”, “float”), empty std::string if attribute not found

void **AddOperation**(const std::string &variable, const std::string &operatorType, const Params &parameters = Params())

Adds operation and parameters to current *IO* object

**Parameters**

- **variable** – variable to add operator to
- **operatorType** – operator type to define
- **parameters** – key/value settings particular to the *IO*, not to be confused by op own parameters

std::string **EngineType**() const

Inspect current engine type from SetEngine

**Returns**

current engine type

### 12.1.4 Variable <T> class

```
template<class T>
```

```
class Variable
```

#### Public Functions

**Variable**() = default

Empty (default) constructor, use it as a placeholder for future variables from *IO::DefineVariable*<T> or *IO::InquireVariable*<T>. Can be used with STL containers.

**~Variable**() = default

Default, using RAII STL containers

explicit **operator bool**() const noexcept

Checks if object is valid, e.g. if( variable ) { //..valid }

void **SetMemorySpace**(const MemorySpace mem)

Sets the memory space for all following Puts/Gets to either host (default) or device

#### Parameters

**mem** – memory space where Put/Get buffers are allocated

MemorySpace **GetMemorySpace**()

Get the memory space that was set by the application

#### Returns

the memory space stored in the *Variable* object

void **SetShape**(const adios2::Dims &shape)

Set new shape, care must be taken when reading back the variable for different steps. Only applies to Global arrays.

#### Parameters

**shape** – new shape dimensions array

void **SetBlockSelection**(const size\_t blockID)

Read mode only. Required for reading local variables, *ShapeID()* = ShapeID::LocalArray or ShapeID::LocalValue. For Global Arrays it will Set the appropriate Start and Count Selection for the global array coordinates.

#### Parameters

**blockID** – variable block index defined at write time. Blocks can be inspected with bpls -D variableName

void **SetSelection**(const adios2::Box<adios2::Dims> &selection)

Sets a variable selection modifying current {start, count} Count is the dimension from Start point

#### Parameters

**selection** – input {start, count}

void **SetMemorySelection**(const adios2::Box<adios2::Dims> &memorySelection)

Set the local start (offset) point to the memory pointer passed at Put and the memory local dimensions (count). Used for non-contiguous memory writes and reads (e.g. multidimensional ghost-cells). Currently Get only works for formats based on BP3.

**Parameters****memorySelection** – {memoryStart, memoryCount}void **SetStepSelection**(const adios2::Box<size\_t> &stepSelection)

Sets a step selection modifying current startStep, countStep countStep is the number of steps from startStep point

**Parameters****stepSelection** – input {startStep, countStep}void **SetAccuracy**(const adios2::Accuracy &a)

Sets the requested accuracy for the next read operation. The actual accuracy after the read is provided in [GetAccuracy\(\)](#)

size\_t **SelectionSize**() const

Returns the number of elements required for pre-allocation based on current count and stepsCount

**Returns**

elements of type T required for pre-allocation

std::string **Name**() const

Inspects [Variable](#) name

**Returns**

name

std::string **Type**() const

Inspects [Variable](#) type

**Returns**

type string literal containing the type: double, float, unsigned int, etc.

size\_t **Sizeof**() const

Inspects size of the current element type, sizeof(T)

**Returns**

sizeof(T) for current system

adios2::ShapeID **ShapeID**() const

Inspects shape id for current variable

**Returns**

from enum adios2::ShapeID

adios2::Dims **Shape**(const size\_t step = adios2::EngineCurrentStep) const

Inspects shape in global variables

**Parameters**

**step** – input for a particular Shape if changing over time. If default, either return absolute or in streaming mode it returns the shape for the current engine step

**Returns**

shape vector

adios2::Dims **Start**() const

Inspects current start point

**Returns**

start vector



adios2::Dims **Count**() const

Inspects current count from start

**Returns**

count vector

size\_t **Steps**() const

For readRandomAccess mode, inspect the number of available steps

**Returns**

available steps

size\_t **StepsStart**() const

For readRandomAccess mode, inspect the start step for available steps

**Returns**

available start step

size\_t **BlockID**() const

For read mode, retrieve current BlockID, default = 0 if not set with SetBlockID

**Returns**

current block id

size\_t **AddOperation**(const *Operator* op, const adios2::Params &parameters = adios2::Params())

Adds operation and parameters to current *Variable* object

**Parameters**

- **op** – operator to be added
- **parameters** – key/value settings particular to the *Variable*, not to be confused by op own parameters

**Returns**

operation index handler in *Operations*()

std::vector<*Operator*> **Operations**() const

Inspects current operators added with AddOperator

**Returns**

vector of Variable<T>::OperatorInfo

void **RemoveOperations**()

Removes all current Operations associated with AddOperation. Provides the possibility to apply or not operators on a block basis.

std::pair<*T*, *T*> **MinMax**(const size\_t step = adios2::DefaultSizeT) const

Read mode only: return minimum and maximum values for current variable at a step. For streaming mode (BeginStep/EndStep): use default (leave empty) for current *Engine* Step At random access mode (File Engines only): default = absolute MinMax

**Parameters**

**step** – input step

**Returns**

pair.first = min pair.second = max

*T* **Min**(const size\_t step = adios2::DefaultSizeT) const

Read mode only: return minimum values for current variable at a step. For streaming mode (within BeginStep/EndStep): use default (leave empty) for current *Engine* Step At random access mode (File Engines only): default = absolute MinMax

**Parameters****step** – input step**Returns**

variable minimum

**T Max**(const size\_t step = adios2::DefaultSizeT) const

Read mode only: return minimum values for current variable at a step. For streaming mode (within BeginStep/EndStep): use default (leave empty) for current *Engine* Step At random access mode (File Engines only): default = absolute MinMax

**Parameters****step** – input step**Returns**

variable minimum

adios2::Accuracy **GetAccuracy**()

Get the provided accuracy for the last read operation. Most operations provide data as it was written, meaning that error is reported as 0.0

std::vector<std::vector<typename *Variable*<T>::Info>> **AllStepsBlocksInfo**()

Read mode only and random-access (no BeginStep/EndStep) with file engines only. Allows inspection of variable info on a per relative step (returned vector index) basis

**Returns**

first vector: relative steps, second vector: blocks info within a step

struct **Info**

Contains block information for a particular Variable<T>

**Public Functions**const *T* \***Data**() const

reference to internal block data (used by inline *Engine*). For deferred variables, valid pointer is not returned until EndStep/PerformGets has been called.

**Public Members**adios2::Dims **Start**

block start

adios2::Dims **Count**

block count

IOType **Min** = IOType()

block Min, if IsValue is false

IOType **Max** = IOType()

block Max, if IsValue is false

`IOType` **Value** = `IOType()`  
block Value, if `IsValue` is true

`int` **WriterID** = 0  
WriterID, source for stream ID that produced this block

`size_t` **BlockID** = 0  
blockID for Block Selection

`size_t` **Step** = 0  
block corresponding step

`bool` **IsReverseDims** = false  
true: Dims were swapped from column-major, false: not swapped

`bool` **IsValue** = false  
true: value, false: array

### 12.1.5 Attribute <T> class

`template<class T>`

`class` **Attribute**

#### Public Functions

**Attribute()** = default

Empty (default) constructor, use it as a placeholder for future attributes from [IO::DefineAttribute<T>](#) or [IO::InquireAttribute<T>](#). Can be used with STL containers.

explicit **operator bool()** const noexcept

Checks if object is valid, e.g. `if( attribute ) { //..valid }`

`std::string` **Name()** const

Inspect attribute name

#### Returns

unique name identifier

`std::string` **Type()** const

Inspect attribute type

#### Returns

type

`std::vector<T>` **Data()** const

Inspect attribute data

#### Returns

data

bool **IsValue()** const

Distinguish single-value attributes from vector attributes

**Returns**

true if single-value, false otherwise

## 12.1.6 Engine class

class **Engine**

### Public Functions

**Engine()** = default

Empty (default) constructor, use it as a placeholder for future engines from *IO::Open*. Can be used with STL containers.

**~Engine()** = default

Using RAII STL containers only

explicit **operator bool()** const noexcept

true: valid engine, false: invalid, not created with *IO::Open* or post *IO::Close*

std::string **Name()** const

Inspect engine name

**Returns**

name from *IO::Open*

std::string **Type()** const

From ADIOS2 engine type: “bpfile”, “sst”, “dataman”, “insitumpi”, “hdf5”

**Returns**

engine type as lower case string

Mode **OpenMode()** const

Returns the Mode used at Open for current *Engine*

**Returns**

StepStatus **BeginStep()**

Begin a logical adios2 step, overloaded version with timeoutSeconds = 0 and mode = Read Check each engine documentation for MPI collective/non-collective behavior.

**Returns**

current step status

StepStatus **BeginStep**(const StepMode mode, const float timeoutSeconds = -1.f)

Begin a logical adios2 step, overloaded version for advanced stream control Check each engine documentation for MPI collective/non-collective behavior.

**Parameters**

- **mode** – see enum `adios2::StepMode` for options, Read is the common use case
- **timeoutSeconds** –

**Returns**

current step status

size\_t **CurrentStep**() const

Inspect current logical step

**Returns**

current logical step

template&lt;class T&gt;

*Variable*<T>::Span **Put**(*Variable*<T> variable, const bool initialize, const T &value)

Put signature that provides access to the internal engine buffer for a pre-allocated variable including a fill value. Returns a fixed size Span (based on C++20 std::span) so applications can populate data value after this Put and before PerformPuts/EndStep. Requires a call to PerformPuts, EndStep, or Close to extract the Min/Max bounds.

**Parameters**

- **variable** – input variable
- **initialize** – bool flag indicating if allocated memory should be initialized with the provided value. Some engines (BP3/BP4) may initialize the allocated memory anyway to zero if this flag is false.
- **value** – provide an initial fill value

**Returns**

span to variable data in engine internal buffer

template&lt;class T&gt;

*Variable*<T>::Span **Put**(*Variable*<T> variable)

Put signature that provides access to an internal engine buffer (decided by the engine) for a pre-allocated variable. Allocated buffer may or may not be initialized to zero by the engine (e.g. BP3/BP4 does, BP5 does not).

**Parameters****variable** – input variable**Returns**

span to variable data in engine internal buffer

template&lt;class T&gt;

void **Put**(*Variable*<T> variable, const T \*data, const Mode launch = Mode::Deferred)Put data associated with a *Variable* in the *Engine***Parameters**

- **variable** – contains variable metadata information
- **data** – user data to be associated with a variable
- **launch** – mode policy

**Throws**

std::invalid\_argument – for invalid variable or nullptr data

template&lt;class T&gt;

void **Put**(const std::string &variableName, const T \*data, const Mode launch = Mode::Deferred)Put data associated with a *Variable* in the *Engine* Overloaded version that accepts a variable name string.**Parameters**

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open

- **data** – user data to be associated with a variable
- **launch** – mode policy

**Throws**

`std::invalid_argument` – if variable not found or nullptr data

template<class T>

void **Put**(*Variable*<T> variable, const T &datum, const Mode launch = Mode::Deferred)

Put data associated with a *Variable* in the *Engine* Overloaded version that accepts r-values and single variable data.

**Parameters**

- **variable** – contains variable metadata information
- **datum** – user data to be associated with a variable, r-value or single data value
- **launch** – mode policy, optional for API consistency, internally is always sync

**Throws**

`std::invalid_argument` – if variable is invalid or nullptr &datum

template<class T>

void **Put**(const std::string &variableName, const T &datum, const Mode launch = Mode::Deferred)

Put data associated with a *Variable* in the *Engine* Overloaded version that accepts variables names, and r-values and single variable data.

**Parameters**

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open
- **datum** – user data to be associated with a variable r-value or single data value
- **launch** – mode policy, optional for API consistency, internally is always sync

**Throws**

`std::invalid_argument` – if variable is invalid or nullptr &datum

template<class T, typename U, class = typename std::enable\_if<std::is\_convertible<U, AdiosView<U>>::value>::type>

inline void **Put**(*Variable*<T> variable, U const &data, const Mode launch = Mode::Deferred)

The next two Put functions are used to accept a variable, and an AdiosViews which is a placeholder for Kokkos::View

**Parameters**

- **variable** – contains variable metadata information
- **data** – represents any user defined object that is not a vector (used for an AdiosView)
- **launch** – mode policy, optional for API consistency, internally is always sync

void **PerformPuts**()

Perform all Put calls in Deferred mode up to this point. Specifically, this causes Deferred data to be copied into *ADIOS* internal buffers as if the Put had been done in Sync mode.

void **PerformDataWrite**()

Write already-Put() array data to disk. If supported by the engine, this may relieve memory pressure by clearing *ADIOS* buffers. It is a collective call and can only be called between Begin/EndStep pairs.

template<class T>

void **Get**(*Variable*<*T*> variable, *T* \*data, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*

#### Parameters

- **variable** – contains variable metadata information
- **data** – user data to be associated with a variable, it must be pre-allocated
- **launch** – mode policy

#### Throws

std::invalid\_argument – for invalid variable or nullptr data

template<class T>

void **Get**(const std::string &variableName, *T* \*data, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*. Overloaded version to get variable by name.

#### Parameters

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open
- **data** – user data to be associated with a variable. It must be pre-allocated
- **launch** – mode policy

#### Throws

std::invalid\_argument – for invalid variableName (variable doesn't exist in *IO*) or nullptr data

template<class T>

void **Get**(*Variable*<*T*> variable, *T* &datum, const Mode launch = Mode::Deferred)

Get single value data associated with a *Variable* from the *Engine* Overloaded version that accepts r-values and single variable data.

#### Parameters

- **variable** – contains variable metadata information
- **datum** – user data to be populated, r-value or single data value
- **launch** – mode policy, optional for API consistency, internally is always sync

#### Throws

std::invalid\_argument – if variable is invalid or nullptr &datum

template<class T>

void **Get**(const std::string &variableName, *T* &datum, const Mode launch = Mode::Deferred)

Get single value data associated with a *Variable* from the *Engine* Overloaded version that accepts r-values and single variable data.

#### Parameters

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open
- **datum** – user data to be populated, r-value or single data value
- **launch** – mode policy, optional for API consistency, internally is always sync

#### Throws

std::invalid\_argument – for invalid variableName (variable doesn't exist in *IO*) or nullptr data

template<class T>

void **Get**(*Variable*<*T*> variable, std::vector<*T*> &dataV, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*. Overloaded version that accepts a std::vector without pre-allocation.

**Parameters**

- **variable** – contains variable metadata information
- **dataV** – user data vector to be associated with a variable, it doesn't need to be pre-allocated. *Engine* will resize.
- **launch** – mode policy

**Throws**

std::invalid\_argument – for invalid variable

template<class *T*>

void **Get**(const std::string &variableName, std::vector<*T*> &dataV, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*. Overloaded version that accepts a std::vector without pre-allocation.

**Parameters**

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open or BeginStep (streaming mode)
- **dataV** – user data vector to be associated with a variable, it doesn't need to be pre-allocated. *Engine* will resize.
- **launch** – mode policy

**Throws**

std::invalid\_argument – for invalid variableName (variable doesn't exist in *IO*)

template<class *T*>

void **Get**(*Variable*<*T*> variable, typename *Variable*<*T*>::Info &info, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*. Data is associated with a block selection, and data is retrieved from variable's BlockInfo.

---

**Note:** Preliminary, experimental API, may change soon.

---

**Parameters**

- **variable** – contains variable metadata information
- **info** – block info struct associated with block selection, call will link with implementation's block info.
- **launch** – mode policy

**Throws**

std::invalid\_argument – for invalid variable or nullptr data

template<class *T*>

void **Get**(const std::string &variableName, typename *Variable*<*T*>::Info &info, const Mode launch = Mode::Deferred)

Get data associated with a *Variable* from the *Engine*. Data is associated with a block selection, and data is retrieved from variable's BlockInfo. Overloaded version to get variable by name.



---

**Note:** Preliminary, experimental API, may change soon.

---

### Parameters

- **variableName** – find variable by name inside *IO* that created this *Engine* with Open or BeginStep (streaming mode)
- **info** – block info struct associated with block selection, call will link with implementation's block info.
- **launch** – mode policy

### Throws

`std::invalid_argument` – for invalid variableName (variable doesn't exist in *IO*)

```
template<class T>
```

```
void Get(Variable<T> variable, T **data) const
```

Assign the value of data to the start of the internal *ADIOS* buffer for variable variable. The value is immediately available.

```
template<class T, typename U, class = typename std::enable_if<std::is_convertible<U,
```

```
AdiosView<U>>::value>::type>
```

```
inline void Get(Variable<T> variable, U const &data, const Mode launch = Mode::Deferred)
```

The next two Get functions are used to accept a variable, and an AdiosViews which is a placeholder for Kokkos::View

### Parameters

- **variable** – contains variable metadata information
- **data** – represents any user defined object that is not a vector (used for an AdiosView)
- **launch** – mode policy, optional for API consistency, internally is always sync

```
void PerformGets()
```

Perform all Get calls in Deferred mode up to this point

```
void EndStep()
```

Ends current step, by default calls PerformsPut/Get internally Check each engine documentation for MPI collective/non-collective behavior.

```
bool BetweenStepPairs()
```

Returns True if engine status is between *BeginStep()*/EndStep() pair, False otherwise.

```
void Flush(const int transportIndex = -1)
```

Manually flush to underlying transport to guarantee data is moved

### Parameters

**transportIndex** –

```
void Close(const int transportIndex = -1)
```

Closes current engine, after this call an engine becomes invalid MPI Collective, calls MPI\_Comm\_free for duplicated communicator at Open

### Parameters

**transportIndex** –

```
template<class T>
```

```
std::map<size_t, std::vector<typename Variable<T>::Info>> AllStepsBlocksInfo(const Variable<T>  
variable) const
```

Extracts all available blocks information for a particular variable. This can be an expensive function, memory scales up with metadata: steps and blocks per step Valid in read mode only.

**Parameters**

**variable** –

**Returns**

map with all variable blocks information

```
template<class T>
```

```
std::vector<typename Variable<T>::Info> BlocksInfo(const Variable<T> variable, const size_t step) const
```

Extracts all available blocks information for a particular variable and step. Valid in read mode only.

**Parameters**

- **variable** – input variable
- **step** – input from which block information is extracted

**Returns**

vector of blocks with info for each block per step, if step not found it returns an empty vector

```
template<class T>
```

```
std::vector<size_t> GetAbsoluteSteps(const Variable<T> variable) const
```

Get the absolute steps of a variable in a file. This is for information purposes only, because absolute steps cannot be used in any ADIOS2 calls.

```
size_t Steps() const
```

Inspect total number of available steps, use for file engines in read mode only

**Returns**

available steps in engine

```
void LockWriterDefinitions()
```

Promise that no more definitions or changes to defined variables will occur. Useful information if called before the first *EndStep()* of an output *Engine*, as it will know that the definitions are complete and constant for the entire lifetime of the output and may optimize metadata handling.

```
void LockReaderSelections()
```

Promise that the reader data selections of are fixed and will not change in future timesteps. This information, provided before the *EndStep()* representing a fixed read pattern, may be utilized by the input *Engine* to optimize data flow.

## 12.1.7 Operator class

```
class Operator
```

## Public Functions

**Operator()** = default

Empty (default) constructor, use it as a placeholder for future operators from *ADIOS::DefineOperator* functions. Can be used with STL containers.

explicit **operator bool()** const noexcept

true: valid object, false: invalid object

std::string **Type()** const noexcept

Inspect current *Operator* type

### Returns

type as string, if invalid returns an empty std::string

void **SetParameter**(const std::string key, const std::string value)

Set a key/value parameters associated with this operator (global parameter from the object it's applied to: *Variable*, *IO*). If key exists, it replace the current value.

### Parameters

- **key** – parameter key
- **value** – parameter value

Params &**Parameters()** const

Inspect current operator parameters

### Returns

map of key/value parameters

## Debugging

For debugging, ADIOS2 C++11 class instances and enums can be passed directly to ostream, as well as converted to human-readable strings via the ubiquitous ToString(object) member variable. You can also directly pass objects to an ostream.

Example:

```
auto myVar = io.DefineVariable<double>("myVar");
std::cout << myVar << " has shape id " << myVar.ShapeID() << std::endl;

// will print:
// Variable<double>(Name: "myVar") has shape id ShapeID::GlobalValue

if (myVar.ShapeID() != adios2::ShapeID::GlobalArray)
{
    throw std::invalid_argument("can't handle " +
                                ToString(myVar.ShapeID()) + " in " +
                                ToString(myVar));
}

// will throw exception like this:
// C++ exception with description "can't handle ShapeID::GlobalValue
// in Variable<double>(Name: "myVar")" thrown
```

## Group API

class **Group**

## Public Functions

`std::vector<std::string> AvailableGroups()`

returns available groups on the path set

### Param

`std::vector<std::string> AvailableVariables()`

returns available variables on the path set

### Param

`std::vector<std::string> AvailableAttributes()`

returns available attributes on the path set

### Param

`std::string InquirePath()`

returns the current path

### Param

`void setPath(std::string path)`

set the path, points to a particular node on the tree

### Parameters

**next** – possible path extension

*Group* `InquireGroup(std::string group_name)`

returns a new group object

### Parameters

**name** – of the group

### Returns

new group object

`template<class T>`

*Variable*`<T> InquireVariable(const std::string &name)`

Gets an existing variable of primitive type by name. A wrapper for the corresponding function of the *IO* class.

### Parameters

**name** – of variable to be retrieved

### Returns

pointer to an existing variable in current *IO*, nullptr if not found

`template<class T>`

*Attribute*`<T> InquireAttribute(const std::string &name, const std::string &variableName = "", const std::string separator = "/")`

Gets an existing attribute of primitive type by name. A wrapper for the corresponding function of the *IO* class

### Parameters

**name** – of attribute to be retrieved

### Returns

pointer to an existing attribute in current *IO*, nullptr if not found

DataType **VariableType**(const std::string &name) const

Inspects variable type. This function can be used in conjunction with MACROS in an else if (type == adios2::GetType<T>() ) {} loop

**Parameters**

**name** – unique variable name identifier in current *IO*

**Returns**

type as in adios2::GetType<T>() (e.g. “double”, “float”), empty std::string if variable not found

DataType **AttributeType**(const std::string &name) const

Inspects attribute type. This function can be used in conjunction with MACROS in an else if (type == adios2::GetType<T>() ) {} loop

**Parameters**

**name** – unique attribute name identifier in current *IO*

**Returns**

type as in adios2::GetType<T>() (e.g. “double”, “float”), empty std::string if attribute not found

The Group API can be used for inquiring other group objects, variables, and attributes by reading, provided that variable and attribute names were written in a tree-like way similar that is used in a file system:

```
"group1/group2/group3/variable1"
"group1/group2/group3/attribute1"
```

A group object containing the tree structure is obtained by the InquireGroup function of the IO object.

```
Group g = io.InquireGroup("group1");
```

Another group object can be generated by the predecessor group object.

```
Group g1 = g.InquireGroup("group2");
```

The absolute path can be inquired or set explicitly

```
std::string path = g.InquirePath();
g.setPath("group1/group2/group3");
```

Names of available groups, variables and attributes could be inquired:

```
std::vector<std::string> groups = g.AvailableGroups();
std::vector<std::string> variables = g.AvailableVariables();
std::vector<std::string> attributes = g.AvailableVariables();
```

Finally, variables can be inquired

```
auto var = g.InquireVariable("variable1");
```

An extra function is provided that returns a type of a variable

```
DataType varType = g.VariableType("variable1");
```

## Step selection

Steps for reading can be selected using a pre-set parameter with as a file name as a key and a list of selected steps separated by comma as a value.

```
io.SetParameter(filename, "1,3");
```

## 12.2 Fortran bindings

The Fortran API is a collection of subroutine calls. The first argument is usually a Fortran type (struct) to an ADIOS2 component, while the last argument is an error integer flag, **integer** `ierr`. `ierr==0` represents successful execution whereas a non-zero value represents an error or a different state. ADIOS2 Fortran bindings provide a list of possible errors coming from the C++ standardized error exception library:

```
! error possible values for ierr
integer, parameter :: adios2_error_none = 0
integer, parameter :: adios2_error_invalid_argument = 1,
integer, parameter :: adios2_error_system_error = 2,
integer, parameter :: adios2_error_runtime_error = 3,
integer, parameter :: adios2_error_exception = 4
```

Click here for a [Fortran write and read example](#) to illustrate the use of the APIs calls. This test will compile under your `build/bin/` directory.

The following subsections describe the overall components and subroutines in the Fortran bindings API.

### 12.2.1 ADIOS2 typed handlers

ADIOS2 Fortran bindings handlers are mapped 1-to-1 to the ADIOS2 components described in the [Components Overview](#) section. For convenience, each type handler contains descriptive components used for read-only inspection.

```
type(adios2_adios)      :: adios
type(adios2_io)         :: io
type(adios2_variable)   :: variable
type(adios2_attribute)  :: attribute
type(adios2_engine)     :: engine

!Read-only components for inspection and ( = defaults)

type adios2_adios
  logical :: valid = .false.
end type

type adios2_io
  logical :: valid = .false.
  character(len=15) :: engine_type = 'BPFile'
end type

type adios2_variable
  logical :: valid = .false.
  character(len=4095) :: name = ''
  integer :: type = -1
  integer :: ndims = -1
end type
```

(continues on next page)

(continued from previous page)

```

type adios2_attribute
  logical :: valid = .false.
  character(len=4095):: name = ''
  integer :: type = -1
  integer :: length = 0
end type

type adios2_engine
  logical :: valid = .false.
  character(len=63):: name = ''
  character(len=15):: type = ''
  integer :: mode = adios2_mode_undefined
end type

type adios2_operator
  logical :: valid = .false.
  character(len=63):: name = ''
  character(len=63):: type = ''
end type

```

**Caution:** Use the type read-only components for information purposes only. Changing their values directly, *e.g.* `variable%name = new_name` does not have any effect inside the ADIOS2 library

## 12.2.2 ADIOS subroutines

- **subroutine** `adios2_init` starting point for the ADIOS2 library

```

! MPI versions
subroutine adios2_init(adios, comm, ierr)
subroutine adios2_init(adios, config_file, comm, ierr)

! Non-MPI serial versions
subroutine adios2_init(adios, ierr)
subroutine adios2_init(adios, config_file, ierr)

! WHERE:

! ADIOS2 handler to allocate
type(adios2_adios), intent(out):: adios

! MPI Communicator
integer, intent(in):: comm

! Optional runtime configuration file (*.xml), see Runtime Configuration
! Files
character*(*), intent(in) :: config_file

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_declare_io` spawn/create an IO component

```
subroutine adios2_declare_io(io, adios, io_name, ierr)

! WHERE:

! output ADIOS2 IO handler
type(adios2_io), intent(out):: io

! ADIOS2 component from adios2_init spawning io tasks
type(adios2_adios), intent(in):: adios

! unique name associated with this IO component inside ADIOS2
character*(*), intent(in):: io_name

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_at_io` retrieve an existing io component. Useful when the original IO handler goes out of scope

```
subroutine adios2_at_io(io, adios, io_name, ierr)

! WHERE:

! output IO handler
type(adios2_io), intent(out):: io

! ADIOS2 component from adios2_init that owns IO tasks
type(adios2_adios), intent(in):: adios

! unique name associated with an existing IO component (created with adios2_
↳ declare_io)
character*(*), intent(in):: io_name

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_define_operator` define an ADIOS2 data compression/reduction operator

```
subroutine adios2_define_operator(op, adios, op_name, op_type, ierr)

! WHERE

! Operator handler
type(adios2_operator), intent(out) :: op

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! Operator name
character*(*), intent(in) :: op_name

! Operator type
```

(continues on next page)



(continued from previous page)

```

character*(*), intent(in)  :: op_type

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_inquire_operator` inquire an ADIOS2 data compression/reduction operator

```

subroutine adios2_inquire_operator(op, adios, op_name, ierr)

! WHERE

! Operator handler
type(adios2_operator), intent(out) :: op

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! Operator name
character*(*), intent(in)  :: op_name

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_flush_all` flush all current engines in all IO objects

```

subroutine adios2_flush_all(adios, ierr)

! WHERE:

! ADIOS2 component from adios2_init owning IO objects and engines
type(adios2_adios), intent(in):: adios

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_remove_io` DANGER ZONE: remove an IO object. This will effectively eliminate any parameter from the config xml file

```

subroutine adios2_remove_io(result, adios, name, ierr)

! WHERE

! Returns True if IO was found, False otherwise
logical, intent(out):: result

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! IO input name
character*(*), intent(in):: name

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_remove_all_ios` DANGER ZONE: remove all IO objects created for this ADIOS2 handler. This will effectively eliminate any parameter from the config xml file as well.

```
subroutine adios2_remove_all_ios(adios, ierr)

! WHERE

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_finalize` final point for the ADIOS2 component

```
subroutine adios2_finalize(adios, ierr)

! WHERE:

! ADIOS2 handler to be deallocated
type(adios2_adios), intent(in):: adios

! error code
integer, intent(out) :: ierr
```

**Caution:** Make sure that for every call to `adios2_init` there is a call to `adios2_finalize` for the same ADIOS2 handler. Not doing so will result in memory leaks.

- **subroutine** `adios2_enter_computation_block` inform ADIOS2 about entering communication-free computation block in main thread. Useful when using Async IO.

```
subroutine adios2_enter_computation_block(adios, ierr)

! WHERE

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_exit_computation_block` inform ADIOS2 about exiting communication-free computation block in main thread. Useful when using Async IO.

```
subroutine adios2_exit_computation_block(adios, ierr)

! WHERE

! ADIOS2 handler
type(adios2_adios), intent(in) :: adios

! error code
integer, intent(out) :: ierr
```

### 12.2.3 IO subroutines

- **subroutine** `adios2_set_engine` set the engine type, see *Supported Engines* for a list of available engines

```

subroutine adios2_set_engine(io, engine_type, ierr)

  ! WHERE:

  ! IO component
  type(adios2_io), intent(in):: io

  ! engine_type: BP (default), HDF5, DataMan, SST, SSC
  character*(*), intent(in):: engine_type

  ! error code
  integer, intent(out) :: ierr

```

- **subroutine** `adios2_in_config_file` checks if an IO object exists in a config file passed to ADIOS2.

```

subroutine adios2_in_config_file(result, io, ierr)

  ! WHERE

  ! Output result to indicate whether IO exists
  logical, intent(out):: result

  ! IO handler
  type(adios2_io), intent(in):: io

  ! error code
  integer, intent(out):: ierr

```

- **subroutine** `adios2_set_parameter` set IO key/value pair parameter in an IO object, see *Supported Engines* for a list of available parameters for each engine type

```

subroutine adios2_set_parameter(io, key, value, ierr)

  ! WHERE:

  ! IO component owning the attribute
  type(adios2_io), intent(in):: io

  ! key in the key/value pair parameter
  character*(*), intent(in):: key

  ! value in the key/value pair parameter
  character*(*), intent(in):: value

  ! error code
  integer, intent(out) :: ierr

```

- **subroutine** `adios2_set_parameters` set a map of key/value parameters in an IO object. Replaces any existing parameters. Otherwise use `set_parameter` for adding single parameters.

```
subroutine adios2_set_parameters(io, parameters, ierr)

! WHERE

! IO handler
type(adios2_io), intent(in) :: io

! Comma-separated parameter list. E.g. "Threads=2, CollectiveMetadata=OFF"
character*(*), intent(in) :: parameters

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_get_parameter` get parameter value from IO object for a given parameter name

```
subroutine adios2_get_parameter(value, io, key, ierr)

! WHERE

! parameter value
character(len=:), allocatable, intent(out) :: value

! IO handler
type(adios2_io), intent(in) :: io

! parameter key to look for in the IO object
character*(*), intent(in) :: key

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_clear_parameters` clear all parameters from the IO object

```
subroutine adios2_clear_parameters(io, ierr)

! WHERE

! IO handler
type(adios2_io), intent(in) :: io

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_add_transport` add a transport to current IO. Must be supported by the currently used engine.

```
subroutine adios2_add_transport(transport_index, io, type, ierr)

! WHERE

! returns a transport_index handler
integer, intent(out):: transport_index
```

(continues on next page)

(continued from previous page)

```

! IO handler
type(adios2_io), intent(in) :: io

! transport type. must be supported by the engine. CAN'T use the keywords
↳ "Transport" or "transport"
character*(*), intent(in) :: type

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_set_transport_parameter` set a parameter for a transport. Overwrites existing parameter with the same key.

```

subroutine adios2_set_transport_parameter(io, transport_index, key, value,
↳ ierr)

! WHERE

! IO handler
type(adios2_io), intent(in):: io

! transport_index handler
integer, intent(in):: transport_index

! transport key
character*(*), intent(in) :: key

! transport value
character*(*), intent(in) :: value

! error code
integer, intent(out):: ierr

```

- **subroutine** `adios2_available_variables` get a list of available variables

```

subroutine adios2_available_variables(io, namestruct, ierr)

! WHERE

! IO handler
type(adios2_io), intent(in) :: io

! name struct handler
type(adios2_namestruct), intent(out) :: namestruct

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_retrieve_names` retrieve variable names from `namestruct` obtained from `adios2_available_variables`. `namelist` must be pre-allocated.

```

subroutine adios2_retrieve_names(namestruct, namelist, ierr)

```

(continues on next page)

(continued from previous page)

```

! WHERE

! namestruct obtained from adios2_available_variables
type(adios2_namestruct), intent(inout) :: namestruct

! namelist that will contain variable names
character(*), dimension(*), intent(inout) :: namelist

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_available_attributes` get list of attributes in the IO object

```

subroutine adios2_available_attributes(io, namestruct, ierr)

! WHERE

! IO handler
type(adios2_io), intent(in) :: io

! list of available attributes
type(adios2_namestruct), intent(out) :: namestruct

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_flush_all_engines` flush all existing engines opened by this IO object

```

subroutine adios2_flush_all_engines(io, ierr)

! WHERE:

! IO in which search and flush for all engines is performed
type(adios2_io), intent(in) :: io

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_remove_variable` remove an existing variable from an IO object

```

subroutine adios2_remove_variable(io, name, result, ierr)

! WHERE:

! IO in which search and removal for variable is performed
type(adios2_io), intent(in) :: io

! unique key name to search for variable
character*(*), intent(in) :: name

! true: variable removed, false: variable not found, not removed
logical, intent(out) :: result

```

(continues on next page)

(continued from previous page)

```
! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_remove_all_variables` remove all existing variables from an IO object

```
subroutine adios2_remove_all_variables(io, ierr)

! WHERE:

! IO in which search and removal for all variables is performed
type(adios2_io), intent(in) :: io

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_remove_attribute` remove existing attribute by its unique name

```
subroutine adios2_remove_attribute(io, name, result, ierr)

! WHERE:

! IO in which search and removal for attribute is performed
type(adios2_io), intent(in) :: io

! unique key name to search for attribute
character*(*), intent(in) :: name

! true: attribute removed, false: attribute not found, not removed
logical, intent(out) :: result

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_remove_all_attributes` remove all existing attributes

```
subroutine adios2_remove_all_attributes(io, ierr)

! WHERE:

! IO in which search and removal for all attributes is performed
type(adios2_io), intent(in) :: io

! error code
integer, intent(out) :: ierr
```

## 12.2.4 Variable subroutines

- **subroutine** `adios2_define_variable` define/create a new variable

```
! Global array variables
subroutine adios2_define_variable(variable, io, variable_name, adios2_type,
↪ &
                                ndims, shape_dims, start_dims, count_dims,
↪ &
                                adios2_constant_dims, ierr)

! Global single value variables
subroutine adios2_define_variable(variable, io, variable_name, adios2_type,
↪ ierr)

! WHERE:

! handler to newly defined variable
type(adios2_variable), intent(out):: variable

! IO component owning the variable
type(adios2_io), intent(in):: io

! unique variable identifier within io
character*(*), intent(in):: variable_name

! defines variable type from adios2 parameters, see next
integer, intent(in):: adios2_type

! number of dimensions
integer, value, intent(in):: ndims

! variable shape, global size, dimensions
! to create local variables optional pass adios2_null_dims
integer(kind=8), dimension(:), intent(in):: shape_dims

! variable start, local offset, dimensions
! to create local variables optional pass adios2_null_dims
integer(kind=8), dimension(:), intent(in):: start_dims

! variable count, local size, dimensions
integer(kind=8), dimension(:), intent(in):: count_dims

! error code
integer, intent(out) :: ierr

! .true. : constant dimensions, shape, start and count won't change
!         (mesh sizes, number of nodes)
!         adios2_constant_dims = .true. use for code clarity
! .false. : variable dimensions, shape, start and count could change
!         (number of particles)
!         adios2_variable_dims = .false. use for code clarity
logical, value, intent(in):: adios2_constant_dims
```

- available `adios2_type` parameters in **subroutine** `adios2_define_variable`



```

integer, parameter :: adios2_type_character = 0
integer, parameter :: adios2_type_real = 2
integer, parameter :: adios2_type_dp = 3
integer, parameter :: adios2_type_complex = 4
integer, parameter :: adios2_type_complex_dp = 5

integer, parameter :: adios2_type_integer1 = 6
integer, parameter :: adios2_type_integer2 = 7
integer, parameter :: adios2_type_integer4 = 8
integer, parameter :: adios2_type_integer8 = 9

integer, parameter :: adios2_type_string = 10
integer, parameter :: adios2_type_string_array = 11

```

**Tip:** Always prefer using `adios2_type_xxx` parameters explicitly rather than raw numbers. *e.g.* use `adios2_type_dp` instead of 3

- **subroutine** `adios2_inquire_variable` inquire and get a variable. See *variable%valid* to check if variable exists.

```

subroutine adios2_inquire_variable(variable, io, name, ierr)

! WHERE:

! output variable handler:
! variable%valid = .true. points to valid found variable
! variable%valid = .false. variable not found
type(adios2_variable), intent(out) :: variable

! IO in which search for variable is performed
type(adios2_io), intent(in) :: io

! unique key name to search for variable
character*(*), intent(in) :: name

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_set_shape` set new `shape_dims` for a variable if its dims are marked as varying in the define call `adios2_define_variable`

```

subroutine adios2_set_shape(variable, ndims, shape_dims, ierr)

! WHERE

! variable handler
type(adios2_variable), intent(in) :: variable

! number of dimensions in shape_dims
integer, intent(in) :: ndims

! new shape_dims

```

(continues on next page)

(continued from previous page)

```
integer(kind=8), dimension(:), intent(in):: shape_dims

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_set_selection` selects part of a variable through `start_dims` and `count_dims`

```
subroutine adios2_set_selection(variable, ndims, start_dims, count_dims,
->ierr)

! WHERE

! variable handler
type(adios2_variable), intent(in) :: variable

! number of dimensions in start_dims and count_dims
integer, intent(in) :: ndims

! new start_dims
integer(kind=8), dimension(:), intent(in):: start_dims

! new count_dims
integer(kind=8), dimension(:), intent(in):: count_dims

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_set_step_selection` set a list of steps by specifying the starting step and the step count

```
subroutine adios2_set_step_selection(variable, step_start, step_count, ierr)

! WHERE

! variable handler
type(adios2_variable), intent(in) :: variable

! new step_start
integer(kind=8), intent(in):: step_start

! new step_count (or number of steps to read from step_start)
integer(kind=8), intent(in):: step_count

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_variable_max` get the maximum value in the variable array

```
subroutine adios2_variable_max(maximum, variable, ierr)

! WHERE

! scalar variable that will contain the maximum value
Generic Fortran types, intent(out) :: maximum
```

(continues on next page)

(continued from previous page)

```

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_variable_min` get the minimum value in the variable array

```

subroutine adios2_variable_min(minimum, variable, ierr)

! WHERE

! scalar variable that will contain the minimum value
Generic Fortran types, intent(out) :: minimum

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_add_operation` add an operation to a variable

```

subroutine adios2_add_operation(operation_index, variable, op, key, value,
↳ierr)

! WHERE

! reference to the operator handle that will be created
integer, intent(out):: operation_index

! variable handler
type(adios2_variable), intent(in):: variable

! Operator handler
type(adios2_operator), intent(in):: op

! Operator key
character*(*), intent(in):: key

! Operator value
character*(*), intent(in):: value

! error code
integer, intent(out):: ierr

```

- **subroutine** `adios2_set_operation_parameter` set a parameter for a operator. Replaces value if parameter already exists.

```

subroutine adios2_set_operation_parameter(variable, operation_index, key,
↳value, ierr)

```

(continues on next page)

(continued from previous page)

```

! WHERE

! variable handler
type(adios2_variable), intent(in):: variable

! Operation index handler
integer, intent(in):: operation_index

! parameter key
character*(*), intent(in):: key

! parameter value
character*(*), intent(in):: value

! error code
integer, intent(out):: ierr

```

- **subroutine** `adios2_variable_name` retrieve variable name

```

subroutine adios2_variable_name(name, variable, ierr)

! WHERE

! variable name
character(len=:), allocatable, intent(out) :: name

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_variable_type` retrieve variable datatype

```

subroutine adios2_variable_type(type, variable, ierr)

! WHERE

! variable type
integer, intent(out) :: type

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_variable_ndims` retrieve number of dimensions for a variable

```

subroutine adios2_variable_ndims(ndims, variable, ierr)

! WHERE

```

(continues on next page)

(continued from previous page)

```

! No. of dimensions
integer, intent(out) :: ndims

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_variable_shape` retrieve the shape of a variable

```

subroutine adios2_variable_shape(shape_dims, ndims, variable, ierr)

! WHERE

! array that contains the shape
integer(kind=8), dimension(:), allocatable, intent(out) :: shape_dims

! no. of dimensions
integer, intent(out) :: ndims

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_variable_steps` retrieve the number of available steps

```

subroutine adios2_variable_steps(steps, variable, ierr)

! WHERE

! no. of steps
integer(kind=8), intent(out) :: steps

! variable handler
type(adios2_variable), intent(in) :: variable

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_set_block_selection` Read mode only. Required for reading local variables. For global arrays it will set the appropriate Start and Count selection for the global array coordinates.

```

subroutine adios2_set_block_selection(variable, block_id, ierr)

! WHERE

! variable handler
type(adios2_variable), intent(in) :: variable

! variable block index defined at write time. Blocks can be inspected with

```

(continues on next page)

(continued from previous page)

```
↪`bpls -D variableName`  
integer(kind=8), intent(in) :: block_id  
  
! error code  
integer, intent(out) :: ierr
```

- **subroutine** `adios2_set_memory_selection` set the local start (offset) point to the memory pointer passed at `adios2_put` and the memory local dimensions (count). Used for non-contiguous memory writes and reads (e.g. multidimensional ghost-cells). Currently Get only works for formats based on BP.

```
subroutine adios2_set_memory_selection(variable, ndims, memory_start_dims, ↪  
↪memory_count_dims, ierr)  
  
! WHERE  
  
! variable handler  
type(adios2_variable), intent(in) :: variable  
  
! no. of dimensions of the variable  
integer, intent(in) :: ndims  
  
! memory start offsets  
integer(kind=8), dimension(:), intent(in) :: memory_start_dims  
  
! no. of elements in each dimension  
integer(kind=8), dimension(:), intent(in) :: memory_count_dims  
  
! error code  
integer, intent(out) :: ierr
```

- **subroutine** `adios2_set_step_selection` set a step selection modifying current `step_start`, `step_count`. `step_count` is the number of steps from `step_start`

```
subroutine adios2_set_step_selection(variable, step_start, step_count, ierr)  
  
! WHERE  
  
! variable handler  
type(adios2_variable), intent(in) :: variable  
  
! starting step  
integer(kind=8), intent(in) :: step_start  
  
! no. of steps from start  
integer(kind=8), intent(in) :: step_count  
  
! error code  
integer, intent(out) :: ierr
```

- **subroutine** `adios2_remove_operations` remove all current operations associated with the variable. Provides the possibility to apply operators on a block basis.

```

subroutine adios2_remove_operations(variable, ierr)

  ! WHERE

  ! variable handler
  type(adios2_variable), intent(in):: variable

  ! error code
  integer, intent(out):: ierr

```

## 12.2.5 Engine subroutines

- **subroutine** `adios2_open` opens an engine to execute IO tasks

```

  ! MPI version: duplicates communicator from adios2_init
  ! Non-MPI serial version
  subroutine adios2_open(engine, io, name, adios2_mode, ierr)

  ! MPI version only to pass a communicator other than the one from adios_init
  subroutine adios2_open(engine, io, name, adios2_mode, comm, ierr)

  ! WHERE:

  ! handler to newly opened adios2 engine
  type(adios2_engine), intent(out) :: engine

  ! IO that spawns an engine based on its configuration
  type(adios2_io), intent(in) :: io

  ! unique engine identifier within io, file name for default BPFile engine
  character*(*), intent(in) :: name

  ! Optional MPI communicator, only in MPI library
  integer, intent(in) :: comm

  ! error code
  integer, intent(out) :: ierr

  ! open mode parameter:
  !
  !         adios2_mode_write,
  !         adios2_mode_append,
  !         adios2_mode_read,
  integer, intent(in):: adios2_mode

```

- **subroutine** `adios2_begin_step` begin a new step or progress to the next step. Starts from 0

```

subroutine adios2_begin_step(engine, adios2_step_mode, timeout_seconds, ↵
↵status, ierr)
  ! Default Timeout = -1.    (block until step available)
  subroutine adios2_begin_step(engine, adios2_step_mode, ierr)
  ! Default step_mode for read and write
  subroutine adios2_begin_step(engine, ierr)

```

(continues on next page)

(continued from previous page)

```

! WHERE

! engine handler
type(adios2_engine), intent(in) :: engine

! step_mode parameter:
!   adios2_step_mode_read (read mode default)
!   adios2_step_mode_append (write mode default)
integer, intent(in):: adios2_step_mode

! optional
! engine timeout (if supported), in seconds
real, intent(in):: timeout_seconds

! status of the stream from adios2_step_status_* parameters
integer, intent(out):: status

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_current_step` extracts current step number

```

subroutine adios2_current_step(current_step, engine, ierr)

! WHERE:

! engine handler
type(adios2_engine), intent(in) :: engine

! populated with current_step value
integer(kind=8), intent(out) :: current_step

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_steps` inspect total number of available steps, use for file engines in read mode only

```

subroutine adios2_steps(steps, engine, ierr)

! WHERE:

! engine handler
type(adios2_engine), intent(in) :: engine

! populated with steps value
integer(kind=8), intent(out) :: steps

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_end_step` end current step and execute transport IO (flush or read).



```

subroutine adios2_end_step(engine, ierr)

  ! WHERE:

  ! engine handler
  type(adios2_engine), intent(in) :: engine

  ! error code
  integer, intent(out) :: ierr

```

- **subroutine** `adios2_put` put variable data and metadata into adios2 for IO operations. Default is deferred mode. For optional sync mode, see *Put: modes and memory contracts*. Variable and data types must match.

```

subroutine adios2_put(engine, variable, data, adios2_mode, ierr)

  ! Default adios2_mode_deferred
  subroutine adios2_put(engine, variable, data, ierr)

  ! WHERE:

  ! engine handler
  type(adios2_engine), intent(in) :: engine

  ! variable handler containing metadata information
  type(adios2_variable), intent(in) :: variable

  ! Fortran bindings supports data types from adios2_type in variables,
  ! up to 6 dimensions
  ! Generic Fortran type from adios2_type
  Generic Fortran types, intent(in):: data
  Generic Fortran types, dimension(:), intent(in):: data
  Generic Fortran types, dimension(,:), intent(in):: data
  Generic Fortran types, dimension(:,:), intent(in):: data
  Generic Fortran types, dimension(:,:,:), intent(in):: data
  Generic Fortran types, dimension(:,:,:,:), intent(in):: data
  Generic Fortran types, dimension(:,:,:,:), intent(in):: data
  Generic Fortran types, dimension(:,:,:,:), intent(in):: data

  ! mode:
  ! adios2_mode_deferred: won't execute until adios2_end_step, adios2_perform_
  ↪ puts or adios2_close
  ! adios2_mode_sync: special case, put data immediately, can be reused after_
  ↪ this call
  integer, intent(in):: adios2_mode

  ! error code
  integer, intent(out) :: ierr

```

- **subroutine** `adios2_perform_puts` execute deferred calls to `adios2_put`

```

subroutine adios2_perform_puts(engine, ierr)

  ! WHERE:

```

(continues on next page)

(continued from previous page)

```

! engine handler
type(adios2_engine), intent(in) :: engine

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_get` get variable data into ADIOS2 for IO operations. Default is deferred mode. For optional sync mode, see *Get: modes and memory contracts*. Variable and data types must match, variable can be obtained from `adios2_inquire_variable`. Memory for data must be pre-allocated.

```

subroutine adios2_get(engine, variable, data, adios2_mode, ierr)

! Default adios2_mode_deferred
subroutine adios2_get(engine, variable, data, ierr)

! WHERE:

! engine handler
type(adios2_engine), intent(in) :: engine

! variable handler containing metadata information
type(adios2_variable), intent(in) :: variable

! Fortran bindings supports data types from adios2_type in variables,
! up to 6 dimensions. Must be pre-allocated
! Generic Fortran type from adios2_type
Generic Fortran types, intent(out):: data
Generic Fortran types, dimension(:), intent(out):: data
Generic Fortran types, dimension(:,), intent(out):: data
Generic Fortran types, dimension(:,,:), intent(out):: data
Generic Fortran types, dimension(:,,,:), intent(out):: data
Generic Fortran types, dimension(:,:,:,:), intent(out):: data
Generic Fortran types, dimension(:,:,:,:,:), intent(out):: data

! mode:
! adios2_mode_deferred: won't execute until adios2_end_step, adios2_perform_
! ↪ gets or adios2_close
! adios2_mode_sync: special case, get data immediately, can be reused after_
! ↪ this call
integer, intent(in):: adios2_mode

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_perform_gets` execute deferred calls to `adios2_get`

```

subroutine adios2_perform_gets(engine, ierr)

! WHERE:

! engine handler
type(adios2_engine), intent(in) :: engine

```

(continues on next page)

(continued from previous page)

```
! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_close` close engine. May re-open.

```
subroutine adios2_close(engine, ierr)

! WHERE:

! engine handler
type(adios2_engine), intent(in) :: engine

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_io_engine_type` get current engine type

```
subroutine adios2_io_engine_type(type, io, ierr)

! WHERE

! engine type (BP, SST, SSC, HDF5, DataMan)
character(len=:), allocatable, intent(out) :: type

! IO handler
type(adios2_io), intent(in) :: io

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_lock_writer_definitions` promise that no more definitions or changes to defined variables will occur. Useful information if called before the first `EndStep()` of an output Engine, as it will know that the definitions are complete and constant for the entire lifetime of the output and may optimize metadata handling.

```
subroutine adios2_lock_writer_definitions(engine, ierr)

! WHERE

! adios2 engine handler
type(adios2_engine), intent(in) :: engine

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_lock_reader_selections` promise that the reader data selections of are fixed and will not change in future timesteps. This information, provided before the `end_step()` representing a fixed read pattern, may be utilized by the input Engine to optimize data flow.

```
subroutine adios2_lock_reader_selections(engine, ierr)

! WHERE
```

(continues on next page)

(continued from previous page)

```

! adios2 engine handler
  type(adios2_engine), intent(in) :: engine

! error code
  integer, intent(out) :: ierr

```

## 12.2.6 Operator subroutines

- **subroutine** `adios2_operator_type` get current Operator type

```

subroutine adios2_operator_type(type, op, ierr)

! WHERE

! Operator type name. See list of supported operator types.
character(len=:), allocatable, intent(out) :: type

! Operator handler
type(adios2_operator), intent(in) :: op

! error code
integer, intent(out) :: ierr

```

## 12.2.7 Attribute subroutines

- **subroutine** `adios2_define_attribute` define/create a new user attribute

```

! Single value attributes
subroutine adios2_define_attribute(attribute, io, attribute_name, data,
↳ierr)

! 1D array attributes
subroutine adios2_define_attribute(attribute, io, attribute_name, data,
↳elements, ierr)

! WHERE:

! handler to newly defined attribute
type(adios2_attribute), intent(out):: attribute

! IO component owning the attribute
type(adios2_io), intent(in):: io

! unique attribute identifier within io
character*(*), intent(in):: attribute_name

! overloaded subroutine allows for multiple attribute data types
! they can be single values or 1D arrays
Generic Fortran types, intent(in):: data

```

(continues on next page)

(continued from previous page)

```

Generic Fortran types, dimension(:), intent(in):: data

! number of elements if passing a 1D array in data argument
integer, intent(in):: elements

! error code
integer, intent(out) :: ierr

```

- **subroutine** `adios2_inquire_attribute` inquire for existing attribute by its unique name

```

subroutine adios2_inquire_attribute(attribute, io, name, ierr)

! WHERE:

! output attribute handler:
! attribute%valid = .true. points to valid found attribute
! attribute%valid = .false. attribute not found
type(adios2_attribute), intent(out) :: attribute

! IO in which search for attribute is performed
type(adios2_io), intent(in) :: io

! unique key name to search for attribute
character*(*), intent(in) :: name

! error code
integer, intent(out) :: ierr

```

**Caution:** Use the `adios2_remove_*` subroutines with extreme CAUTION. They create outdated dangling information in the `adios2_type` handlers. If you don't need them, don't use them.

- **subroutine** `adios2_attribute_data` retrieve attribute data

```

subroutine adios2_attribute_data(data, attribute, ierr)

! WHERE

! data handler
character*(*), intent(out):: data
real, intent(out):: data
real(kind=8), intent(out):: data
integer(kind=1), intent(out):: data
integer(kind=2), intent(out):: data
integer(kind=4), intent(out):: data
integer(kind=8), intent(out):: data
character*(*), dimension(:), intent(out):: data
real, dimension(:), intent(out):: data
real(kind=8), dimension(:), intent(out):: data
integer(kind=2), dimension(:), intent(out):: data
integer(kind=4), dimension(:), intent(out):: data
integer(kind=8), dimension(:), intent(out):: data

```

(continues on next page)

(continued from previous page)

```
! attribute
type(adios2_attribute), intent(in):: attribute

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_attribute_name` inspect attribute name

```
subroutine adios2_attribute_name(name, attribute, ierr)

! WHERE

! name to be output
character(len=:), allocatable, intent(out) :: name

! attribute handler
type(adios2_attribute), intent(in) :: attribute

! error code
integer, intent(out) :: ierr
```

- **subroutine** `adios2_inquire_variable_attribute` retrieve a handler to a previously defined attribute associated to a variable

```
subroutine adios2_inquire_variable_attribute(attribute, io, attribute_name,   
↵variable_name, separator, ierr)

! WHERE

! attribute handler
type(adios2_attribute), intent(out) :: attribute

! IO handler
type(adios2_io), intent(in) :: io

! attribute name
character*(*), intent(in) :: attribute_name

! variable name
character*(*), intent(in) :: variable_name

! hierarchy separator (e.g. "/" in variable/attribute )
character*(*), intent(in) :: separator

! error code
integer, intent(out) :: ierr
```

## 12.3 C bindings

The C bindings are specifically designed for C applications and those using an older C++ standard (98 and 03). If you are using a C++11 or more recent standard, please use the C++11 bindings.

The C bindings are based on opaque pointers to the components described in the *Components Overview* section.

C API handlers:

```
adios2_adios*
adios2_io*
adios2_variable*
adios2_attribute*
adios2_engine*
adios2_operator*
```

Every ADIOS2 function that generates a new `adios2_*` unique handler returns the latter explicitly. Therefore, checks can be applied to know if the resulting handler is NULL. Other functions used to manipulate these valid handlers will return a value of type `enum adios2_error` explicitly. These possible errors are based on the [C++ standardized exceptions](#). Each error will issue a more detailed description in the standard error output: `stderr`.

`adios2_error` possible values:

```
typedef enum {
    /** success */
    adios2_error_none = 0,

    /**
     * user input error
     */
    adios2_error_invalid_argument = 1,

    /** low-level system error, e.g. system IO error */
    adios2_error_system_error = 2,

    /** runtime errors other than system errors, e.g. memory overflow */
    adios2_error_runtime_error = 3,

    /** any other error exception */
    adios2_error_exception = 4
} adios2_error;
```

Usage:

```
adios2_variable* var = adios2_define_variable(io, ...)
if(var == NULL )
{
    // ... something went wrong with adios2_define_variable
    // ... check stderr
}
else
{
    adios2_type type;
    adios2_error errio = adios2_variable_type(&type, var)
    if(errio){
```

(continues on next page)

(continued from previous page)

```
// ... something went wrong with adios2_variable_type
if( errio == adios2_error_invalid_argument)
{
    // ... user input error
    // ... check stderr
}
}
```

---

**Note:** Use `#include "adios2_c.h"` for the C bindings, `adios2.h` is the C++11 header.

---

### 12.3.1 adios2\_adios handler functions

#### Defines

**adios2\_init**(comm)

**adios2\_init\_config**(config\_file, comm)

#### Functions

adios2\_adios \***adios2\_init\_mpi**(MPI\_Comm comm)

Starting point for MPI apps. Creates an ADIOS handler. MPI collective and it calls MPI\_Comm\_dup

##### Parameters

**comm** – defines domain scope from application

##### Returns

success: handler, failure: NULL

adios2\_adios \***adios2\_init\_config\_mpi**(const char \*config\_file, MPI\_Comm comm)

Starting point for MPI apps. Creates an ADIOS handler allowing a runtime config file. MPI collective and it calls MPI\_Comm\_dup and MPI\_Bcast to pass the configFile contents

##### Parameters

- **config\_file** – runtime configuration file in xml format
- **comm** – defines domain scope from application

##### Returns

success: handler, failure: NULL

adios2\_adios \***adios2\_init\_serial**(void)

Initialize an ADIOS struct pointer handler in a serial, non-MPI application. Doesn't require a runtime config file.

##### Returns

success: handler, failure: NULL

adios2\_adios \***adios2\_init\_config\_serial**(const char \*config\_file)

Initialize an ADIOS struct pointer handler in a serial, non-MPI application. Doesn't require a runtime config file.



**Returns**

success: handler, failure: NULL

adios2\_io \***adios2\_declare\_io**(adios2\_adios \*adios, const char \*name)

Declares a new io handler

**Parameters**

- **adios** – owner the io handler
- **name** – unique io identifier within current adios handler

**Returns**

success: handler, failure: NULL

adios2\_io \***adios2\_declare\_io\_order**(adios2\_adios \*adios, const char \*name, adios2\_arrayordering order)

Declares a new io handler with specific array ordering

**Parameters**

- **adios** – owner the io handler
- **name** – unique io identifier within current adios handler
- **order** – array ordering

**Returns**

success: handler, failure: NULL

adios2\_io \***adios2\_at\_io**(adios2\_adios \*adios, const char \*name)

Retrieves a previously declared io handler by name

**Parameters**

- **adios** – owner the io handler
- **name** – unique name for the previously declared io handler

**Returns**

success: handler, failure: NULL

adios2\_operator \***adios2\_define\_operator**(adios2\_adios \*adios, const char \*name, const char \*type)

Defines an adios2 supported operator by its type.

**Parameters**

- **adios** – owner the op handler
- **name** – unique operator name identifier within current ADIOS object
- **type** – supported ADIOS2 operator type: zfp, sz

**Returns**

success: handler, failure: NULL

adios2\_operator \***adios2\_inquire\_operator**(adios2\_adios \*adios, const char \*name)

Retrieves a previously defined operator handler

**Parameters**

- **adios** – owner the op handler
- **name** – unique name for the previously defined op handler

**Returns**

success: handler, failure: NULL

adios2\_error **adios2\_flush\_all**(adios2\_adios \*adios)

Flushes all adios2\_engine in write mode in all adios2\_io handlers. If no adios2\_io or adios2\_engine exists it does nothing.

**Parameters**

**adios** – owner of all io and engines to be flushed

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_finalize**(adios2\_adios \*adios)

Final point for adios handler. Deallocates adios pointer. Required to avoid memory leaks. MPI collective and it calls MPI\_Comm\_free

**Parameters**

**adios** – handler to be deallocated, must be initialized with adios2\_init or adios2\_init\_config

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_remove\_io**(adios2\_bool \*result, adios2\_adios \*adios, const char \*name)

DANGER ZONE: removes an io created with adios2\_declare\_io. Will create dangling pointers for all the handlers inside removed io. NOTE: Use result, not adios2\_error to check if the IO was removed.

**Parameters**

- **result** – output adios2\_true: io not found and not removed, adios2\_false: io not found and not removed
- **adios** – owner of io to be removed
- **name** – input unique identifier for io to be removed

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_remove\_all\_ios**(adios2\_adios \*adios)

DANGER ZONE: removes all ios created with adios2\_declare\_io. Will create dangling pointers for all the handlers inside all removed io.

**Parameters**

**adios** – owner of all ios to be removed

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_enter\_computation\_block**(adios2\_adios \*adios)

Inform ADIOS about entering communication-free computation block in main thread. Useful when using Async IO

adios2\_error **adios2\_exit\_computation\_block**(adios2\_adios \*adios)

Inform ADIOS about exiting communication-free computation block in main thread. Useful when using Async IO

## 12.3.2 adios2\_io handler functions

### Functions

adios2\_error **adios2\_in\_config\_file**(adios2\_bool \*result, const adios2\_io \*io)

Check if io exists in a config file passed to the adios handler that created this io.

#### Parameters

- **result** – output adios2\_true=1: in config file, adios2\_false=0: not in config file
- **io** – handler

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_engine**(adios2\_io \*io, const char \*engine\_type)

Set the engine type for current io handler.

#### Parameters

- **io** – handler
- **engine\_type** – predefined engine type, default is bpfile

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_parameters**(adios2\_io \*io, const char \*parameters)

Set several parameters at once.

#### Parameters

- **io** – handler
- **string** – parameters in the format “param1=value1 , param2 = value2”

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_parameter**(adios2\_io \*io, const char \*key, const char \*value)

Set a single parameter. Overwrites value if key exists.

#### Parameters

- **io** – handler
- **key** – parameter key
- **value** – parameter value

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_get\_parameter**(char \*value, size\_t \*size, const adios2\_io \*io, const char \*key)

Return IO parameter value string and length without ‘\0’ character For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for ‘\0’ if desired), then call the function again with the buffer. Then ‘\0’ terminate it if desired.

#### Parameters

- **value** – output
- **size** – output, value size without ‘\0’

- **io** – input handler
- **key** – input parameter key, if not found return size = 0 and value is untouched

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_clear\_parameters**(adios2\_io \*io)

Clear all parameters.

**Parameters**

**io** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_add\_transport**(size\_t \*transport\_index, adios2\_io \*io, const char \*type)

Add a transport to current io handler. Must be supported by current engine type.

**Parameters**

- **transport\_index** – handler used for setting transport parameters or at adios2\_close
- **io** – handler
- **type** – must be a supported transport type for a particular Engine. CAN'T use the keywords "Transport" or "transport"

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_transport\_parameter**(adios2\_io \*io, const size\_t transport\_index, const char \*key, const char \*value)

Set a single parameter to an existing transport identified with a transport\_index handler from adios2\_add\_transport. Overwrites existing parameter with the same key.

**Parameters**

- **io** – handler
- **transport\_index** – handler from adios2\_add\_transport
- **key** – parameter key
- **value** – parameter value

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_variable \***adios2\_define\_variable**(adios2\_io \*io, const char \*name, const adios2\_type type, const size\_t ndims, const size\_t \*shape, const size\_t \*start, const size\_t \*count, const adios2\_constant\_dims constant\_dims)

Define a variable within io.

**Parameters**

- **io** – handler that owns the variable
- **name** – unique variable identifier
- **type** – primitive type from enum adios2\_type in adios2\_c\_types.h
- **ndims** – number of dimensions
- **shape** – global dimension

- **start** – local offset
- **count** – local dimension
- **constant\_dims** – adios2\_constant\_dims\_true:: shape, start, count won't change;  
adios2\_constant\_dims\_false: shape, start, count will change after definition

#### Returns

success: handler, failure: NULL

adios2\_variable \***adios2\_inquire\_variable**(adios2\_io \*io, const char \*name)

Retrieve a variable handler within current io handler.

#### Parameters

- **io** – handler to variable io owner
- **name** – unique variable identifier within io handler

#### Returns

found: handler, not found: NULL

adios2\_error **adios2\_inquire\_all\_variables**(adios2\_variable \*\*\*variables, size\_t \*size, adios2\_io \*io)

Returns an array of variable handlers for all variable present in the io group

#### Parameters

- **variables** – output array of variable pointers (pointer to an adios2\_variable\*\*)
- **size** – output number of variables
- **io** – handler to variables io owner

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_inquire\_group\_variables**(adios2\_variable \*\*\*variables, const char \*full\_group\_name,  
size\_t \*size, adios2\_io \*io)

adios2\_attribute \***adios2\_define\_attribute**(adios2\_io \*io, const char \*name, const adios2\_type type, const void  
\*value)

Define an attribute value inside io.

#### Parameters

- **io** – handler that owns the attribute
- **name** – unique attribute name inside IO handler
- **type** – primitive type from enum adios2\_type in adios2\_c\_types.h
- **value** – attribute single value

#### Returns

success: handler, failure: NULL

adios2\_attribute \***adios2\_define\_attribute\_array**(adios2\_io \*io, const char \*name, const adios2\_type type,  
const void \*data, const size\_t size)

Define an attribute array inside io.

#### Parameters

- **io** – handler that owns the attribute
- **name** – unique attribute name inside IO handler

- **type** – primitive type from enum `adios2_type` in `adios2_c_types.h`
- **data** – attribute data array
- **size** – number of elements of data array

**Returns**

success: handler, failure: NULL

```
adios2_attribute *adios2_define_variable_attribute(adios2_io *io, const char *name, const adios2_type
                                                    type, const void *value, const char *variable_name,
                                                    const char *separator)
```

Define an attribute single value associated to an existing variable by its name

**Parameters**

- **io** – handler that owns the variable and attribute
- **name** – unique attribute name inside a variable in io handler
- **type** – primitive type from enum `adios2_type` in `adios2_c_types.h`
- **value** – attribute single value
- **variable\_name** – unique variable identifier in io handler. If variable doesn't exist `adios2_error` is `adios2_error_invalid_argument`.
- **separator** – hierarchy separator (e.g. "/" in `variable_name/name` )

**Returns**

success: handler, failure: NULL

```
adios2_attribute *adios2_define_variable_attribute_array(adios2_io *io, const char *name, const
                                                         adios2_type type, const void *data, const size_t
                                                         size, const char *variable_name, const char
                                                         *separator)
```

Define an attribute array associated to an existing variable by its name

**Parameters**

- **io** – handler that owns the variable and attribute
- **name** – unique attribute name inside a variable in io handler
- **type** – primitive type from enum `adios2_type` in `adios2_c_types.h`
- **data** – attribute data single value or array
- **size** – number of elements of data array
- **variable\_name** – unique variable identifier in io handler. If variable doesn't exist `adios2_error` is true.
- **separator** – hierarchy separator (e.g. "/" in `variable/attribute` )

**Returns**

success: handler, failure: NULL

```
adios2_attribute *adios2_inquire_attribute(adios2_io *io, const char *name)
```

Returns a handler to a previously defined attribute by name

**Parameters**

- **io** – handler to attribute io owner
- **name** – unique attribute identifier within io handler

**Returns**

found: handler, not found: NULL

adios2\_attribute \***adios2\_inquire\_variable\_attribute**(adios2\_io \*io, const char \*name, const char \*variable\_name, const char \*separator)

Retrieve a handler to a previously defined attribute associated to a variable

**Parameters**

- **io** – handler to attribute and variable io owner
- **name** – unique attribute name inside a variable in io handler
- **variable\_name** – name of the variable associate with this attribute
- **separator** – hierarchy separator (e.g. “/” in variable/attribute )

**Returns**

found: handler, not found: NULL

adios2\_error **adios2\_inquire\_all\_attributes**(adios2\_attribute \*\*\*attributes, size\_t \*size, adios2\_io \*io)

Returns an array of attribute handlers for all attribute present in the io group

**Parameters**

- **attributes** – output array of attribute pointers (pointer to an adios2\_attribute\*\*)
- **size** – output number of attributes
- **io** – handler to attributes io owner

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_inquire\_group\_attributes**(adios2\_attribute \*\*\*attributes, const char \*full\_prefix, size\_t \*size, adios2\_io \*io)

adios2\_error **adios2\_inquire\_subgroups**(char \*\*\*subGroupNames, const char \*full\_prefix, size\_t \*size, adios2\_io \*io)

Return a list of list sub group names

adios2\_error **adios2\_remove\_variable**(adios2\_bool \*result, adios2\_io \*io, const char \*name)

DANGEROUS! Removes a variable identified by name. Might create dangling pointers.

**Parameters**

- **result** – output adios2\_true(1): found and removed variable, adios2\_false(0): not found, nothing to remove
- **io** – handler variable io owner
- **name** – unique variable name within io handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_remove\_all\_variables**(adios2\_io \*io)

DANGEROUS! Removes all existing variables in current IO object. Might create dangling pointers.

**Parameters**

**io** – handler variables io owner

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

char \*\***adios2\_available\_variables**(adios2\_io \*io, size\_t \*size)

returns an array or c strings for names of available variables Might create dangling pointers

**Parameters**

- **io** – handler variables io owner
- **length** – of array of strings

**Returns**

names of variables as an array of strings

char \*\***adios2\_available\_attributes**(adios2\_io \*io, size\_t \*size)

returns an array or c strings for names of available attributes Might create dangling pointers

**Parameters**

- **io** – handler variables io owner
- **length** – of array of strings

**Returns**

names of variables as an array of strings

adios2\_error **adios2\_remove\_attribute**(adios2\_bool \*result, adios2\_io \*io, const char \*name)

DANGEROUS! Removes an attribute identified by name. Might create dangling pointers.

**Parameters**

- **result** – output adios2\_true(1): found and removed attribute, adios2\_false(0): not found, nothing to remove
- **io** – handler attribute io owner
- **name** – unique attribute name within io handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_remove\_all\_attributes**(adios2\_io \*io)

DANGEROUS! Removes all existing attributes in current IO object. Might create dangling pointers.

**Parameters**

**io** – handler attributes io owner

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_engine \***adios2\_open**(adios2\_io \*io, const char \*name, const adios2\_mode mode)

Open an Engine to start heavy-weight input/output operations. In MPI version reuses the communicator from adios2\_init or adios2\_init\_config MPI Collective function as it calls MPI\_Comm\_dup

**Parameters**

- **io** – engine owner
- **name** – unique engine identifier
- **mode** – adios2\_mode\_write, adios2\_mode\_read, adios2\_mode\_append, adios2\_mode\_readRandomAccess and adios2\_mode\_readFlattenSteps

**Returns**

success: handler, failure: NULL



adios2\_engine \***adios2\_open\_new\_comm**(adios2\_io \*io, const char \*name, const adios2\_mode mode, MPI\_Comm comm)

Open an Engine to start heavy-weight input/output operations. MPI Collective function as it calls MPI\_Comm\_dup

#### Parameters

- **io** – engine owner
- **name** – unique engine identifier
- **mode** – adios2\_mode\_write, adios2\_mode\_read, adios2\_mode\_append, adios2\_mode\_readRandomAccess and adios2\_mode\_readFlattenSteps
- **comm** – communicator other than adios' handler comm. MPI only.

#### Returns

success: handler, failure: NULL

adios2\_error **adios2\_flush\_all\_engines**(adios2\_io \*io)

Flushes all engines created with current io handler using adios2\_open

#### Parameters

**io** – handler whose engine will be flushed

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_engine\_type**(char \*engine\_type, size\_t \*size, const adios2\_io \*io)

return engine type string and length without null character For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

#### Parameters

- **engine\_type** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, engine\_type size without '\0'
- **io** – handler

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_engine \***adios2\_get\_engine**(adios2\_io \*io, const char \*name)

### 12.3.3 adios2\_variable handler functions

#### Functions

adios2\_error **adios2\_set\_shape**(adios2\_variable \*variable, const size\_t ndims, const size\_t \*shape)

Set new shape, care must be taken when reading back the variable for different steps. Only applies to Global arrays.

#### Parameters

- **variable** – handler for which new selection will be applied to
- **ndims** – number of dimensions for start and count
- **shape** – new shape dimensions array

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_memory\_space**(adios2\_variable \*variable, const adios2\_memory\_space mem)

Sets the memory space for all following Puts/Gets to either host (default) or device

**Parameters**

**mem** – memory space where Put/Get buffers are allocated

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_get\_memory\_space**(adios2\_memory\_space \*mem, adios2\_variable \*variable)

Get the memory space that was set by the application for a given variable

**Parameters**

- **memory** – space output, the variable memory space
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_block\_selection**(adios2\_variable \*variable, const size\_t block\_id)

Read mode only. Required for reading local variables. For Global Arrays it will Set the appropriate Start and Count Selection for the global array coordinates.

**Parameters**

- **variable** – handler for which new selection will be applied to
- **block\_id** – variable block index defined at write time. Blocks can be inspected with bpls -D variableName

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_selection**(adios2\_variable \*variable, const size\_t ndims, const size\_t \*start, const size\_t \*count)

Set new start and count dimensions

**Parameters**

- **variable** – handler for which new selection will be applied to
- **ndims** – number of dimensions for start and count
- **start** – new start dimensions array
- **count** – new count dimensions array

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_memory\_selection**(adios2\_variable \*variable, const size\_t ndims, const size\_t \*memory\_start, const size\_t \*memory\_count)

Set the local start (offset) point to the memory pointer passed at Put and the memory local dimensions (count). Used for non-contiguous memory writes and reads (e.g. multidimensional ghost-cells). Currently not working for calls to Get.

**Parameters**

- **variable** – handler for which new memory selection will be applied to

- **ndims** – number of dimensions for `memory_start` and `memory_count`
- **memory\_start** – relative local offset of `variable.start` to the contiguous memory pointer passed at `Put` from which data starts. e.g. if `variable.start = {rank*Ny,0}` and there is 1 ghost cell per dimension, then `memory_start = {1,1}`
- **memory\_count** – local dimensions for the contiguous memory pointer passed at `adios2_put`, e.g. if there is 1 ghost cell per dimension and `variable.count = {Ny,Nx}`, then `memory_count = {Ny+2,Nx+2}`

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_set\_step\_selection**(`adios2_variable *variable`, `const size_t step_start`, `const size_t step_count`)

Set new step selection using `step_start` and `step_count`. Used mostly for reading from file-based engines (e.g. `bpfile`, `hdf5`)

**Parameters**

- **variable** – handler for which new selection will be applied to
- **step\_start** – starting step for reading
- **step\_count** – number of steps to read from step start

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_variable\_name**(`char *name`, `size_t *size`, `const adios2_variable *variable`)

Retrieve variable name For safe use, call this function first with `NULL` name parameter to get the size, then preallocate the buffer (with room for `'\0'` if desired), then call the function again with the buffer. Then `'\0'` terminate it if desired.

**Parameters**

- **name** – output, string without trailing `'\0'`, `NULL` or preallocated buffer
- **size** – output, name size without `'\0'`
- **variable** – handler

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_variable\_type**(`adios2_type *type`, `const adios2_variable *variable`)

Retrieve variable type

**Parameters**

- **type** – output, from enum `adios2_type`
- **variable** – handler

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_variable\_type\_string**(`char *type`, `size_t *size`, `const adios2_variable *variable`)

Retrieve variable type in string form “char”, “unsigned long”, etc. For safe use, call this function first with `NULL` name parameter to get the size, then preallocate the buffer (with room for `'\0'` if desired), then call the function again with the buffer. Then `'\0'` terminate it if desired.

**Parameters**

- **type** – output, string without trailing ‘\0’, NULL or preallocated buffer
- **size** – output, type size without ‘\0’
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_shapeid**(adios2\_shapeid \*shapeid, const adios2\_variable \*variable)

Retrieve variable shapeid

**Parameters**

- **shapeid** – output, from enum adios2\_shapeid
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_ndims**(size\_t \*ndims, const adios2\_variable \*variable)

Retrieve current variable number of dimensions

**Parameters**

- **ndims** – output
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_shape**(size\_t \*shape, const adios2\_variable \*variable)

Retrieve current variable shape

**Parameters**

- **shape** – output, must be pre-allocated with ndims
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_shape\_with\_memory\_space**(size\_t \*shape, const adios2\_variable \*variable,  
const adios2\_memory\_space mem)

Retrieve current variable shape for a given memory space

**Parameters**

- **shape** – output, must be pre-allocated with ndims
- **variable** – handler
- **memory** – space

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_start**(size\_t \*start, const adios2\_variable \*variable)

Retrieve current variable start

**Parameters**

- **start** – output, single value

- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_count**(size\_t \*count, const adios2\_variable \*variable)

Retrieve current variable start

**Parameters**

- **count** – output, must be pre-allocated with ndims
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_steps\_start**(size\_t \*steps\_start, const adios2\_variable \*variable)

Read API, get available steps start from available steps count (e.g. in a file for a variable).

**Parameters**

- **steps\_start** – output absolute first available step, don't use with adios2\_set\_step\_selection as inputs are relative, use 0 instead.
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_steps**(size\_t \*steps, const adios2\_variable \*variable)

Read API, get available steps count from available steps count (e.g. in a file for a variable). Not necessarily contiguous.

**Parameters**

- **steps** – output available steps, single value
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_selection\_size**(size\_t \*size, const adios2\_variable \*variable)

Returns the minimum required allocation (in number of elements of a certain type, not bytes) for the current selection

**Parameters**

- **size** – number of elements of current type to be allocated by a pointer/vector to read current selection
- **variable** – handler for which data size will be inspected from

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_add\_operation**(size\_t \*operation\_index, adios2\_variable \*variable, adios2\_operator \*op, const char \*key, const char \*value)

Adds an operation to a variable (e.g. compression)

**Parameters**

- **operation\_index** – output handler to be used with adios2\_add\_operation\_param

- **variable** – handler on which operation is applied to
- **op** – handler to adios2\_operator associated to current operation
- **key** – parameter key supported by the operation, empty if none
- **value** – parameter value supported by the operation, empty if none

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_set\_operation\_parameter**(adios2\_variable \*variable, const size\_t operation\_id, const char \*key, const char \*value)

Adds a parameter to an operation created with adios2\_add\_operation

**Parameters**

- **variable** – handler on which operation is applied to
- **operation\_id** – handler returned from adios2\_add\_operation
- **key** – parameter key supported by the operation
- **value** – parameter value supported by the operation

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_remove\_operations**(adios2\_variable \*variable)

Removes all current Operations associated with AddOperation. Provides the possibility to apply or not operators on a block basis.

**Parameters**

**variable** – handler on which operation is applied to

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_min**(void \*min, const adios2\_variable \*variable)

Read mode only: return the absolute minimum for current variable

**Parameters**

- **min** – output: variable minimum, must be of the same type as the variable handler
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_variable\_max**(void \*max, const adios2\_variable \*variable)

Read mode only: return the absolute maximum for current variable

**Parameters**

- **max** – output: variable minimum, must be of the same type as the variable handler
- **variable** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

### 12.3.4 adios2\_attribute handler functions

#### Functions

`adios2_error` **adios2\_attribute\_name**(char \*name, size\_t \*size, const adios2\_attribute \*attribute)

Retrieve attribute name For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

##### Parameters

- **name** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, name size without '\0'
- **attribute** – handler

##### Returns

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_attribute\_type**(adios2\_type \*type, const adios2\_attribute \*attribute)

Retrieve attribute type

##### Parameters

- **type** –
- **attribute** – handler

##### Returns

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_attribute\_type\_string**(char \*type, size\_t \*size, const adios2\_attribute \*attribute)

Retrieve attribute type in string form “char”, “unsigned long”, etc. For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

##### Parameters

- **type** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, type size without '\0'
- **attribute** – handler

##### Returns

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_attribute\_is\_value**(adios2\_bool \*result, const adios2\_attribute \*attribute)

Checks if attribute is a single value or an array

##### Parameters

- **result** – output, `adios2_true`: single value, `adios2_false`: array
- **attribute** – handler

##### Returns

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_attribute\_size**(size\_t \*size, const adios2\_attribute \*attribute)

Returns the number of elements (as in C++ STL `size()` function) if attribute is a 1D array. If single value returns 1

**Parameters**

- **size** – output, number of elements in attribute
- **attribute** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_attribute\_data**(void \*data, size\_t \*size, const adios2\_attribute \*attribute)

Retrieve attribute data pointer (read-only)

**Parameters**

- **data** – output attribute values, must be pre-allocated
- **size** – data size
- **attribute** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

### 12.3.5 adios2\_engine handler functions

**Functions**

adios2\_error **adios2\_engine\_name**(char \*name, size\_t \*size, const adios2\_engine \*engine)

Return engine name string and length without '\0' character For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

**Parameters**

- **name** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, engine\_type size without '\0'
- **engine** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_engine\_get\_type**(char \*type, size\_t \*size, const adios2\_engine \*engine)

Return engine type string and length without '\0' character For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

**Parameters**

- **type** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, engine\_type size without '\0'
- **engine** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_engine\_openmode**(adios2\_mode \*mode, const adios2\_engine \*engine)

Return the engine's Open mode.

**Parameters**



- **mode** – output, adios2\_mode parameter used in *adios2\_open()*
- **engine** – handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_begin\_step**(adios2\_engine \*engine, const adios2\_step\_mode mode, const float timeout\_seconds, adios2\_step\_status \*status)

Begin a logical adios2 step stream Check each engine documentation for MPI collective/non-collective behavior.

**Parameters**

- **engine** – handler
- **mode** – see enum adios2\_step\_mode in adios2\_c\_types.h for options, read is the common use case
- **timeout\_seconds** – provide a time out in Engine opened in read mode
- **status** – output from enum adios2\_step\_status in adios2\_c\_types.h

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_current\_step**(size\_t \*current\_step, const adios2\_engine \*engine)

Inspect current logical step

**Parameters**

- **current\_step** – output
- **engine** – input handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_between\_step\_pairs**(size\_t \*between\_step\_pairs, const adios2\_engine \*engine)

Inspect current between step status

**Parameters**

- **between\_step\_pairs** – output boolean
- **engine** – input handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_steps**(size\_t \*steps, const adios2\_engine \*engine)

Inspect total number of available steps, use for file engines in read mode only

**Parameters**

- **steps** – output available steps in engine
- **engine** – input handler

**Returns**

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_put**(adios2\_engine \*engine, adios2\_variable \*variable, const void \*data, const adios2\_mode launch)

Put data associated with a Variable in an engine, used for engines with adios2\_mode\_write at adios2\_open

**Parameters**

- **engine** – handler for a particular engine where data will be put
- **variable** – contains variable metadata information
- **data** – user data to be associated with a variable, must be the same type passed to `adios2_define_variable`
- **launch** – mode launch policy

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_put\_by\_name**(`adios2_engine *engine`, `const char *variable_name`, `const void *data`, `const adios2_mode launch`)

Put data associated with a Variable in an engine, used for engines with `adios2_mode_write` at `adios2_open`. This is the name string version

**Parameters**

- **engine** – handler for a particular engine where data will be put
- **variable\_name** – variable with this name must exists in `adios2_io` that opened the engine handler (1st parameter)
- **data** – user data to be associated with a variable, must be the same type passed to `adios2_define_variable`
- **launch** – mode launch policy

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_perform\_puts**(`adios2_engine *engine`)

Performs all the `adios2_put` and `adios2_put_by_name` called with mode `adios2_mode_deferred`, up to this point, by copying user data into internal ADIOS buffers. User data can be reused after this point.

**Parameters**

**engine** – handler for a particular engine where data will be put

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_perform\_data\_write**(`adios2_engine *engine`)

Write array data to disk. This may relieve memory pressure by clearing ADIOS buffers. It is a collective call. User data can be reused after this point.

**Parameters**

**engine** – handler for a particular engine where data will be put

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_get**(`adios2_engine *engine`, `adios2_variable *variable`, `void *data`, `const adios2_mode launch`)

Gets data associated with a Variable from an engine, used for engines with `adios2_mode_read` at `adios2_open`. This is the name string version

**Parameters**

- **engine** – handler for a particular engine where data will be put
- **variable** – handler must exists in `adios2_io` that opened the engine handler (1st parameter). Typically from `adios2_inquire_variable`

- **data** – user data to be associated with a variable, must be the same type passed to `adios2_define_variable`. Must be pre-allocated for the required variable selection.
- **launch** – mode launch policy

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_get\_by\_name**(`adios2_engine` \*engine, const char \*variable\_name, void \*data, const `adios2_mode` launch)

Gets data associated with a Variable from an engine, used for engines with `adios2_mode_read` at `adios2_open`. This is the name string version

**Parameters**

- **engine** – handler for a particular engine where data will be put
- **variable\_name** – variable with this name must exists in `adios2_io` that opened the engine handler (1st parameter).
- **data** – user data to be associated with a variable, must be the same type passed to `adios2_define_variable`. Must be pre-allocated for the required variable selection.
- **launch** – mode launch policy

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_perform\_gets**(`adios2_engine` \*engine)

Performs all the `adios2_get` and `adios2_get_by_name` called with mode `adios2_mode_deferred` up to this point by getting the data from the Engine. User data can be reused after this point.

**Parameters**

**engine** – handler for a particular engine where data will be obtained

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_end\_step**(`adios2_engine` \*engine)

Terminates interaction with current step. By default puts/gets data to/from all transports Check each engine documentation for MPI collective/non-collective behavior.

**Parameters**

**engine** – handler executing IO tasks

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_flush**(`adios2_engine` \*engine)

Explicit engine buffer flush to transports

**Parameters**

**engine** – input

**Returns**

`adios2_error` 0: success, see enum `adios2_error` for errors

`adios2_error` **adios2\_flush\_by\_index**(`adios2_engine` \*engine, const int transport\_index)

Explicit engine buffer flush to transport index

**Parameters**

- **engine** – input

- **transport\_index** – index to be flushed

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_close**(adios2\_engine \*engine)

Close all transports in adios2\_Engine. Call is required to close system resources. MPI Collective, calls MPI\_Comm\_free for duplicated communicator at Open

#### Parameters

**engine** – handler containing all transports to be closed. NOTE: engines NEVER become NULL after this function is called.

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_close\_by\_index**(adios2\_engine \*engine, const int transport\_index)

Close a particular transport from the index returned by adios2\_add\_transport

#### Parameters

- **engine** – handler containing all transports to be closed. NOTE: engines NEVER become NULL due to this function.
- **transport\_index** – handler from adios2\_add\_transport

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors

adios2\_error **adios2\_lock\_writer\_definitions**(adios2\_engine \*engine)

Promise that no more definitions or changes to defined variables will occur. Useful information if called before the first *adios2\_end\_step()* of an output Engine, as it will know that the definitions are complete and constant for the entire lifetime of the output and may optimize metadata handling.

#### Parameters

**engine** – handler

adios2\_error **adios2\_lock\_reader\_selections**(adios2\_engine \*engine)

Promise that the reader data selections of are fixed and will not change in future timesteps. This information, provided before the EndStep() representing a fixed read pattern, may be utilized by the input Engine to optimize data flow.

#### Parameters

**engine** – handler

adios2\_varinfo \***adios2\_inquire\_blockinfo**(adios2\_engine \*engine, adios2\_variable \*variable, const size\_t step)

Get the list of blocks for a variable in a given step. In Streaming mode, step is unused, always the current step is processed.

#### Returns

Newly allocated adios2\_varinfo structure, NULL pointer if step does not exist. The memory must be freed by the adios2\_free\_blockinfo function

void **adios2\_free\_blockinfo**(adios2\_varinfo \*data\_blocks)

free adios2\_varinfo structure

#### Parameters

**data\_blocks** –

#### Returns

void

### 12.3.6 adios2\_operator handler functions

#### Functions

adios2\_error **adios2\_operator\_type**(char \*type, size\_t \*size, const adios2\_operator \*op)

Retrieve operator type For safe use, call this function first with NULL name parameter to get the size, then preallocate the buffer (with room for '\0' if desired), then call the function again with the buffer. Then '\0' terminate it if desired.

#### Parameters

- **type** – output, string without trailing '\0', NULL or preallocated buffer
- **size** – output, type size without '\0'
- **op** – operator handler to be inspected

#### Returns

adios2\_error 0: success, see enum adios2\_error for errors



## HIGH-LEVEL APIS

The high-level APIs are designed for simple tasks for which performance is not critical. Unlike the *Full APIs*, the high-level APIs only require a single object handler resembling a C++ `fstream` or a Python file I/O idiom. The high-level APIs are recommended to both first-time and advanced users; the low-level APIs being recommended only when performance testing identifies a bottleneck or when more control is needed.

Typical scenarios for using the simple high-level APIs are:

- Reading a file to perform data analysis with libraries (matplotlib, scipy, etc.)
- Interactive: few calls make interactive usage easier.
- Saving data to files is small or personal projects
- Online frameworks: *e.g.* Jupyter notebooks, see python-mpi examples running on [MyBinder](#)

The designed functionality syntax is closely related to the native language IO bindings for formatted text files *e.g.* C++ `fstream` `getline`, and Python file IO. The main function calls are: `open` (or constructor in C++), `write`, `read` and `close` (or destructor in C++). In addition, ADIOS2 borrows the corresponding language native syntax for advancing lines to advance the step in write mode, and for a “step-by-step” streaming basis in read mode. See each language section in this chapter for a write/read example.

---

**Note:** The simplified APIs are based on language native file IO interface. Hence `write` and `read` calls are always synchronized and variables data memory is ready to use immediately after these calls.

---

Currently ADIOS2 support bindings for the following languages and their minimum standards:

Language	Standard	Interface	Based on
C++	11/newer	<code>#include adios2.h</code>	<code>fstream</code>
Matlab			

The following sections provide a summary of the API calls on each language and links to Write and Read examples to put it all together.

## 13.1 C++ High-Level API

C++11 High-Level APIs are based on a single object `adios2::fstream`

**Caution:** DO NOT place use namespace `adios2` in your C++ code. Use `adios2::fstream` directly to prevent conflicts with `std::fstream`.

### 13.1.1 C++11 Write example

```
#include <adios2.h>
...

int rank, size;
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);

// Nx, Ny from application, std::size_t
const adios2::Dims shape{Nx, Ny * static_cast<std::size_t>(size)};
const adios2::Dims start{0, Ny * static_cast<std::size_t>(rank)};
const adios2::Dims count{Nx, Ny};

adios2::fstream oStream("cfld.bp", adios2::fstream::out, MPI_COMM_WORLD);

// NSteps from application
for (std::size_t step = 0; step < NSteps; ++step)
{
    if(rank == 0 && step == 0) // global variable
    {
        oStream.write<int32_t>("size", size);
    }

    // physicalTime double, <double> is optional
    oStream.write<double>("physicalTime", physicalTime );
    // T and P are std::vector<float>
    oStream.write( "temperature", T.data(), shape, start, count );
    // adios2::endl will advance the step after writing pressure
    oStream.write( "pressure", P.data(), shape, start, count, adios2::end_step );
}

// Calling close is mandatory!
oStream.close();
```



### 13.1.2 C++11 Read “step-by-step” example

```
#include <adios2.h>
...

int rank, size;
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);

// Selection Window from application, std::size_t
const adios2::Dims start{0, 0};
const adios2::Dims count{SelX, SelY};

if( rank == 0)
{
    // if only one rank is active use MPI_COMM_SELF
    adios2::fstream iStream("cfid.bp", adios2::fstream::in, MPI_COMM_SELF);

    adios2::fstep iStep;
    while (adios2::getstep(iStream, iStep))
    {
        if( iStep.currentstep() == 0 )
        {
            const std::size_t sizeOriginal = iStep.read<std::size_t>("size");
        }
        const double physicalTime = iStream.read<double>( "physicalTime");
        const std::vector<float> temperature = iStream.read<float>( "temperature", start,
↪count );
        const std::vector<float> pressure = iStream.read<float>( "pressure", start, count,
↪);
    }
    // Don't forget to call close!
    iStream.close();
}
```

### 13.1.3 adios2::fstream API documentation

class **fstream**

#### Public Types

enum **openmode**

Available open modes for *adios2::fstream* constructor or open calls

Values:

enumerator **out**

write

enumerator **in**

read

enumerator **in\_random\_access**

read\_random\_access

enumerator **in\_flatten\_steps**

flatten all input steps to 1

enumerator **app**

append, not yet supported

## Public Functions

**fstream**(const std::string &name, adios2::fstream::openmode mode, MPI\_Comm comm, const std::string engineType = "BPFile")

High-level API MPI constructor, based on C++11 fstream. Allows for passing parameters in source code.

### Parameters

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **comm** – MPI communicator establishing domain for fstream
- **engineType** – available adios2 engine

### Throws

std::invalid\_argument – (user input error) or std::runtime\_error (system error)

**fstream**(const std::string &name, const adios2::fstream::openmode mode, MPI\_Comm comm, const std::string &configFile, const std::string ioInConfigFile)

High-level API MPI constructor, based on C++11 fstream. Allows for runtime config file.

### Parameters

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **comm** – MPI communicator establishing domain for fstream
- **configFile** – adios2 runtime configuration file
- **ioInConfigFile** – specific io name in configFile

### Throws

std::invalid\_argument – (user input error) or std::runtime\_error (system error)

**fstream**(const std::string &name, const adios2::fstream::openmode mode, const std::string engineType = "BPFile")

High-level API non-MPI constructor, based on C++11 fstream. Allows for passing parameters in source code.

### Parameters

- **name** – stream name

- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **engineType** – available adios2 engine

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

**fstream**(const std::string &name, const adios2::fstream::openmode mode, const std::string &configFile, const std::string ioInConfigFile)

High-level API MPI constructor, based on C++11 fstream. Allows for runtime config file.

**Parameters**

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **configFile** – adios2 runtime configuration file
- **ioInConfigFile** – specific io name in configFile

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

**fstream**() = default

Empty constructor, allows the use of open later in the code

**~fstream**() = default

Using RAII STL containers only

explicit **operator bool**() const noexcept

Checks if fstream object is valid

void **open**(const std::string &name, const openmode mode, MPI\_Comm comm, const std::string engineType = "BPFile")

High-level API MPI open, based on C++11 fstream. Allows for passing parameters in source code. Used after empty constructor.

**Parameters**

- **name** – stream name
- **mode** – *adios2::fstream::in* (Read), *adios2::fstream::out* (Write), *adios2::fstream::app* (Append)
- **comm** – MPI communicator establishing domain for fstream
- **engineType** – available adios2 engine

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

void **open**(const std::string &name, const openmode mode, MPI\_Comm comm, const std::string configFile, const std::string ioInConfigFile)

High-level API MPI constructor, based on C++11 fstream. Allows for runtime config file. Used after empty constructor.

**Parameters**

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **comm** – MPI communicator establishing domain for fstream

- **configFile** – adios2 runtime configuration file
- **ioInConfigFile** – specific io name in configFile

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

void **open**(const std::string &name, const *openmode* mode, const std::string engineType = "BPFile")

High-level API non-MPI open, based on C++11 `fstream`. Allows for passing parameters in source code. Used after empty constructor.

**Parameters**

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **engineType** – available adios2 engine

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

void **open**(const std::string &name, const *openmode* mode, const std::string configFile, const std::string ioInConfigFile)

High-level API non-MPI constructor, based on C++11 `fstream`. Allows for runtime config file. Used after empty constructor.

**Parameters**

- **name** – stream name
- **mode** – *fstream::in* (Read), *fstream::out* (Write), *fstream::app* (Append)
- **configFile** – adios2 runtime configuration file
- **ioInConfigFile** – specific io name in configFile

**Throws**

`std::invalid_argument` – (user input error) or `std::runtime_error` (system error)

void **set\_parameter**(const std::string key, const std::string value) noexcept

Set a single stream parameter based on *Engine* supported parameters. MUST be passed before the first call to write or read. See: <https://adios2.readthedocs.io/en/latest/engines/engines.html>

**Parameters**

- **key** – input parameter key
- **value** – input parameter value

void **set\_parameters**(const adios2::Params &parameters) noexcept

Set stream parameters based on *Engine* supported parameters. MUST be passed before the first call to write or read. See: <https://adios2.readthedocs.io/en/latest/engines/engines.html>

**Parameters**

**parameters** – input key/value parameters

template<class T>

void **write\_attribute**(const std::string &name, const T &value, const std::string &variableName = "", const std::string separator = "/", const bool endStep = false)

Define attribute inside `fstream` or for a variable after write. Single value input version.

**Parameters**

- **name** – unique attribute identifier *IO* object or for a *Variable* if variableName is not empty (associated to a variable)
- **value** – single data value
- **variableName** – default is empty, if not empty attributes is associated to a variable after a write
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.
- **endStep** – similar to std::endStep, end current step and flush (default). Use adios2::endStep for true.

```
template<class T>
void write_attribute(const std::string &name, const T *data, const size_t size, const std::string
                    &variableName = "", const std::string separator = "/", const bool endStep = false)
```

Define attribute inside fstream or for a variable after write. Array input version.

#### Parameters

- **name** – unique attribute identifier *IO* object or for a *Variable* if variableName is not empty (associated to a variable)
- **data** – pointer to user data
- **size** – number of data elements
- **variableName** – default is empty, if not empty attributes is associated to a variable after a write
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.
- **endStep** – similar to std::endStep, end current step and flush (default). Use adios2::endStep for true.

```
template<class T>
void write(const std::string &name, const T *data, const adios2::Dims &shape = adios2::Dims(), const
          adios2::Dims &start = adios2::Dims(), const adios2::Dims &count = adios2::Dims(), const bool
          endStep = false)
```

writes a self-describing array variable

#### Parameters

- **name** – variable name
- **data** – variable data data
- **shape** – variable global MPI dimensions. Pass empty for local variables.
- **start** – variable offset for current MPI rank. Pass empty for local variables.
- **count** – variable dimension for current MPI rank. Local variables only have count.
- **endStep** – similar to std::endStep, end current step and flush (default). Use adios2::endStep if true.

#### Throws

std::invalid\_argument – (user input error) or std::runtime\_error (system error)

```
template<class T>
```

```
void write(const std::string &name, const T *data, const adios2::Dims &shape, const adios2::Dims &start,  
          const adios2::Dims &count, const adios2::vParams &operations, const bool endStep = false)
```

write overload that allows passing supported operations (e.g. lossy compression “zfp”, “mgard”, “sz”) to a self-described array variable

#### Parameters

- **name** – variable name
- **data** – variable data data
- **shape** – variable global MPI dimensions. Pass empty for local variables.
- **start** – variable offset for current MPI rank. Pass empty for local variables.
- **count** – variable dimension for current MPI rank. Local variables only have count.
- **operations** – vector of operations, each entry is a std::pair:
- **endStep** – similar to std::endStep, end current step and flush (default). Use adios2::endStep if true.

#### Throws

std::invalid\_argument – (user input error) or std::runtime\_error (system error)

```
template<class T>
```

```
void write(const std::string &name, const T &value, const bool isLocalValue = false, const bool endStep =  
          false)
```

Write a self-describing single-value variable

#### Parameters

- **name** – variable name
- **value** – variable data value (can be r-value)
- **isLocalValue** – true: local value (returned as GlobalArray), false: global value (returned as global value)
- **endStep** – similar to std::endStep, end current step and flush (default). Use adios2::endStep for true.

#### Throws

std::invalid\_argument – (user input error) or std::runtime\_error (system error)

```
template<class T>
```

```
void read(const std::string &name, T *data, const size_t blockID = 0)
```

Reads into a pre-allocated pointer. When used with adios2::getstep reads current step

#### Parameters

- **name** – variable name
- **data** – pre-allocated pointer to hold read data
- **blockID** – required for local variables, specify current block to be selected

#### Throws

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
void read(const std::string &name, T &value, const size_t blockID = 0)
```

Reads a value. When used with adios2::getstep reads current step value

#### Parameters

- **name** – variable name
- **value** – output value, if variable is not found (name and type don't match) the returned value address becomes nullptr
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
void read(const std::string &name, T *data, const size_t stepsStart, const size_t stepsCount = 1, const size_t blockID = 0)
```

Read accessing steps in random access mode. Not be used with `adios2::getstep` as it throw an exception when reading in stepping mode.

**Parameters**

- **name** – variable name
- **data** – pre-allocated pointer to hold read data, if variable is not found (name and type don't match) it becomes nullptr
- **stepsStart** – variable initial step (relative to the variable first appearance, not absolute step in stream)
- **stepsCount** – variable number of steps form `step_start`, don't have to be contiguous, necessarily
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
void read(const std::string &name, T &value, const size_t step, const size_t blockID = 0)
```

Reads into a single value for a single step. Not be used with `adios2::getstep` as it throws an exception when reading in stepping mode.

**Parameters**

- **name** – variable name
- **value** – filled with value, if variable is not found (name, type and step don't match) the returned value address becomes nullptr
- **step** – selected single step
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
void read(const std::string &name, T *data, const adios2::Dims &start, const adios2::Dims &count, const size_t blockID = 0)
```

Reads into a pre-allocated pointer a selection piece in dimension. When used with `adios2::getstep` reads current step

**Parameters**

- **name** – variable name

- **data** – pre-allocated pointer to hold read data, if variable is not found (name and type don't match) it becomes nullptr
- **start** – variable local offset selection
- **count** – variable local dimension selection from start
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
void read(const std::string &name, T *data, const adios2::Dims &start, const adios2::Dims &count, const size_t stepsStart, const size_t stepsCount, const size_t blockID = 0)
```

Reads into a pre-allocated pointer a selection piece in dimensions and steps. Not be used with `adios2::getstep` as it throws an exception when reading in stepping mode.

**Parameters**

- **name** – variable name
- **data** – pre-allocated pointer to hold read data, if variable is not found (name and type don't match) it becomes a nullptr
- **start** – variable local offset selection
- **count** – variable local dimension selection from start
- **stepsStart** – variable initial step (relative to the variable first appearance, not absolute step in stream)
- **stepsCount** – variable number of steps form `step_start`, don't have to be necessarily contiguous
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

```
template<class T>
```

```
std::vector<T> read(const std::string &name, const size_t blockID = 0)
```

Reads entire variable for current step (streaming mode: step by step)

**Parameters**

- **name** – variable name
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

**Returns**

data of variable name for current step. Single data will have a `size=1` vector

```
template<class T>
```

```
std::vector<T> read(const std::string &name, const size_t stepsStart, const size_t stepsCount = 1, const size_t blockID = 0)
```

Returns a vector with full variable dimensions for the current step selection. Not be used with `adios2::getstep` as it throw an exception when reading in stepping mode.

**Parameters**

- **name** – variable name



- **stepsStart** – variable initial step (relative to the variable first appearance, not absolute step in stream)
- **stepsCount** – variable number of steps form step\_start, don't have to be contiguous, necessarily
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

**Returns**

data of variable name for current step, empty if exception is thrown

```
template<class T>
std::vector<T> read(const std::string &name, const Dims &start, const Dims &count, const size_t blockID = 0)
```

Reads a selection piece in dimension for current step (streaming mode: step by step)

**Parameters**

- **name** – variable name
- **start** – variable local offset selection
- **count** – variable local dimension selection from start
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

**Returns**

data of variable name for current step, empty if exception is thrown

```
template<class T>
std::vector<T> read(const std::string &name, const Dims &start, const Dims &count, const size_t stepsStart,
const size_t stepsCount, const size_t blockID = 0)
```

Reads a selection piece in dimension and a selection piece in steps (non-streaming mode). Not be used with adios2::getstep as it throw an exception when reading in stepping mode.

**Parameters**

- **name** – variable name
- **start** – variable local offset selection
- **count** – variable local dimension selection from start
- **stepsStart** – variable initial step (relative to the variable first appearance, not absolute step in stream)
- **stepsCount** – variable number of steps form step\_start, don't have to be contiguous, necessarily
- **blockID** – required for local variables, specify current block to be selected

**Throws**

throws – exception if variable name, dimensions or step not found

**Returns**

variable data, empty if exception is thrown

```
template<class T>
```

```
std::vector<T> read_attribute(const std::string &name, const std::string &variableName = "", const  
                             std::string separator = "/")
```

Reads an attribute returning a vector For single data vector size = 1

#### Parameters

- **name** – attribute name
- **variableName** – default is empty, if not empty look for an attribute associated to a variable
- **separator** – default is “/”, hierarchy between variable name and attribute, e.g. variable-Name/attribute1, variableName::attribute1. Not used if variableName is empty.

#### Returns

vector containing attribute data

```
void end_step()
```

At write: ends the current step At read: use it in streaming mode to inform the writer that the reader is done consuming the step. No effect for file engines.

```
void close()
```

close current stream becoming inaccessible

```
size_t current_step() const noexcept
```

Return current step when getstep is called in a loop, read mode only

#### Returns

current step

### Friends

```
friend bool getstep(adios2::fstream &stream, adios2::fstep &step)
```

Gets step from stream Based on std::getline, enables reading on a step-by-step basis in a while or for loop. Read mode only

#### Parameters

- **stream** – input stream containing steps
- **step** – output object current step, adios2::fstep in an alias to *adios2::fstream* with scope narrowed to one step

#### Returns

true: step is valid, false: step is invalid (end of stream).

## 13.2 Matlab simple bindings

The ADIOS Matlab API supports reading data from ADIOS BP files with a simplified API that consists of three functions:

- ADIOSOPEN returns a structure with information on an ADIOS BP File (variables and attributes).
- ADIOSREAD reads in a variable from the file. It expects the info structure returned by ADIOSOPEN.
- ADIOSCLOSE closes the file.

### 13.2.1 Organization of an ADIOS BP file

An ADIOS BP file contains a set of variables and attributes. Each variable in the group has a path, which defines a logical hierarchy of the variables within the file.

### 13.2.2 Time dimension of a variable

Variables can be written several times from a program, if they have a time dimension. The reader exposes the variables with an extra dimension, i.e. a 2D variable written over time is seen as a 3D variable. In MATLAB, the extra dimension is the last dimension (the slowest changing dimension). Since the reader allows reading an arbitrary slice of a variable, data for one timestep can be read in with slicing.

### 13.2.3 Min/max of arrays

The ADIOS BP format stores the min/max values in each variable. The info structure therefore contains these min/max values. There is practically no overhead to provide this information (along with the values of all attributes) even for file sizes of several terabytes.

In the Matlab console use help for these functions

```
>>> help adiosopen
>>> help adiosread
>>> help adiosclose
```

### 13.2.4 ADIOSOPEN

**FILE = adiosopen(PATH)**

Open a file for reading pointed by PATH and return an information structure (FILE).

The returned FILE structure contains the following information

Name	File path
Handlers	Object handlers to pass on to ADIOS functions
FileHandler	uint64 file handler
GroupHandler	uint64 IO group object handler
ADIOSHandler	uint64 ADIOS object handler
Variables	Structure array of variables
Name	Path of variable
Type	Matlab type class of data
Dims	Array of dimensions
StepsStart	First step's index for this variable in file, always at least 1
StepsCount	Number of steps for this variable in file, always at least 1
GlobalMin	Global minimum of the variable (1-by-1 mxArray)
GlobalMax	Global maximum of the variable
Attribute	Structure array of attributes
Name	Path of attribute

(continues on next page)

(continued from previous page)

Type	Matlab <b>type class</b> of data
Value	Attribute value

### 13.2.5 ADIOSREAD

Read data from a BP file opened with `adiosopen`. Provide the structure returned by `adiosopen` as the first input argument, and the path to a variable. Inspect `file.Variables` and `file.Attributes` for the list of variables and attributes available in a file.

```
data = adiosread(file, VARPATH)
```

Read the entire variable `VARPATH` from a BP file. `file` is the output of `ADIOSOPEN`. `VARPATH` is a string to a variable or attribute. If an N-dimensional array variable has multiple steps in the file this function reads all steps and returns an N+1 dimensional array where the last dimension equals the number of steps.

```
data = adiosread(file, INDEX)
```

Read the entire variable from a BP file. `INDEX` points to a variable in the `file.Variables` array.

```
data = adiosread(..., START, COUNT, STEPSTART, STEPCOUNT)
```

Read a portion of a variable.

#### START and COUNT:

A slice is defined as two arrays of N integers, where N is the number of dimensions of the variable, describing the "start" and "count" values. The "start" values start from 1.

E.g. `[1 5], [10 2]` reads the first 10 values in the first dimension and 2 values from the 5th position in the second dimension resulting in a 10-by-2 array.

You can use negative numbers to index from the **end** of the array as in python. `-1` refers to the last element of the array, `-2` the one before and so on.

E.g. `[-1], [1]` reads in the last value of a 1D array.

`[1], [-1]` reads in the complete 1D array.

#### STEPSTART and STEPCOUNT:

Similarly, the number of steps from a specific step can be read instead of all data. Steps start from 1. Negative index can be used as well.

E.g. `-1, 1` will read in the last step from the file

`n, -1` will read all steps from 'n' to the last one

### 13.2.6 ADIOSCLOSE

```
adiosclose(file)
```

Close file and free internal data structures. `file` is the structure returned by `adiosopen`.

## PYTHON APIS

### 14.1 Python Example Code

The Python APIs follows closely Python style directives. They rely on numpy and, optionally, on mpi4py, if the underlying ADIOS2 library is compiled with MPI.

For online examples on MyBinder :

- [Python-MPI Notebooks](#)
- [Python-noMPI Notebooks](#)

#### 14.1.1 Examples in the ADIOS2 repository

- **Simple file-based examples**
  - `examples/hello/helloWorld/hello-world.py`
  - `examples/hello/bpReader/bpReaderHeatMap2D.py`
  - `examples/hello/bpWriter/bpWriter.py`
- **Staging examples using staging engines SST and DataMan**
  - `examples/hello/sstWriter/sstWriter.py`
  - `examples/hello/sstReader/sstReader.py`
  - `examples/hello/datamanWriter/dataManWriter.py`
  - `examples/hello/datamanReader/dataManReader.py`

#### 14.1.2 Python Write example

```
from mpi4py import MPI
import numpy as np
from adios2 import Stream

comm = MPI.COMM_WORLD
rank = comm.Get_rank()
size = comm.Get_size()

nx = 10
shape = [size * nx]
```

(continues on next page)

(continued from previous page)

```

start = [rank * nx]
count = [nx]

temperature = np.zeros(nx, dtype=np.double)
pressure = np.ones(nx, dtype=np.double)
delta_time = 0.01
physical_time = 0.0
nsteps = 5

with Stream("cfid.bp", "w", comm) as s:
    # NSteps from application
    for _ in s.steps(nsteps):
        if rank == 0 and s.current_step() == 0:
            # write a Python integer
            s.write("nproc", size)

            # write a Python floating point value
            s.write("physical_time", physical_time)
            # temperature and pressure are numpy arrays
            s.write("temperature", temperature, shape, start, count)
            s.write_attribute("temperature/unit", "K")
            s.write("pressure", pressure, shape, start, count)
            s.write_attribute("pressure/unit", "Pa")
            physical_time += delta_time

```

```

$ mpirun -n 4 python3 ./adios2-doc-write.py
$ bpls -la cfid.bp
int64_t  nproc          scalar = 4
double   physical_time  5*scalar = 0 / 0.04
double   pressure       5*{40} = 1 / 1
string   pressure/unit  attr   = "Pa"
double   temperature    5*{40} = 0 / 0
string   temperature/unit attr   = "K"

```

### 14.1.3 Python Read “step-by-step” example

```

import numpy as np
from adios2 import Stream

with Stream("cfid.bp", "r") as s:
    # steps comes from the stream
    for _ in s.steps():

        # track current step
        print(f"Current step is {s.current_step()}")

        # inspect variables in current step
        for name, info in s.available_variables().items():
            print("variable_name: " + name, end=" ")
            for key, value in info.items():

```

(continues on next page)

(continued from previous page)

```

        print("\t" + key + ": " + value, end=" ")
    print()

    if s.current_step() == 0:
        nproc = s.read("nproc")
        print(f"nproc is {nproc} of type {type(nproc)}")

        # read variables return a numpy array with corresponding selection
        physical_time = s.read("physical_time")
        print(f"physical_time is {physical_time} of type {type(physical_time)}")
        temperature = s.read("temperature")
        temp_unit = s.read_attribute("temperature/unit")
        print(f"temperature array size is {temperature.size} of shape {temperature.shape}
↪")
        print(f"temperature unit is {temp_unit} of type {type(temp_unit)}")
        pressure = s.read("pressure")
        press_unit = s.read_attribute("pressure/unit")
        print(f"pressure unit is {press_unit} of type {type(press_unit)}")
    print()

```

```

$ python3 adios2-doc-read.py
Current step is 0
variable_name: nproc      AvailableStepsCount: 1  Max: 4  Min: 4  Shape:      ↪
↪SingleValue: true       Type: int64_t
variable_name: physical_time  AvailableStepsCount: 1  Max: 0  Min: 0  Shape:      ↪
↪SingleValue: true       Type: double
variable_name: pressure      AvailableStepsCount: 1  Max: 1  Min: 1  Shape: 40    ↪
↪SingleValue: false       Type: double
variable_name: temperature   AvailableStepsCount: 1  Max: 0  Min: 0  Shape: 40    ↪
↪SingleValue: false       Type: double
nproc is 4 of type <class 'numpy.ndarray'>
physical_time is 0.0 of type <class 'numpy.ndarray'>
temperature array size is 40 of shape (40,)
temperature unit is K of type <class 'str'>
pressure unit is Pa of type <class 'str'>

Current step is 1
variable_name: physical_time  AvailableStepsCount: 1  Max: 0.01  Min: 0.01  ↪
↪Shape:      SingleValue: true  Type: double
variable_name: pressure      AvailableStepsCount: 1  Max: 1  Min: 1  Shape: 40    ↪
↪SingleValue: false       Type: double
variable_name: temperature   AvailableStepsCount: 1  Max: 0  Min: 0  Shape: 40    ↪
↪SingleValue: false       Type: double
physical_time is 0.01 of type <class 'numpy.ndarray'>
temperature array size is 40 of shape (40,)
temperature unit is K of type <class 'str'>
pressure unit is Pa of type <class 'str'>

...

```

### 14.1.4 Python Read Random Access example

```
import numpy as np
from adios2 import FileReader

with FileReader("cfv.bp") as s:
    # inspect variables
    vars = s.available_variables()
    for name, info in vars.items():
        print("variable_name: " + name, end=" ")
        for key, value in info.items():
            print("\t" + key + ": " + value, end=" ")
        print()
    print()

    nproc = s.read("nproc")
    print(f"nproc is {nproc} of type {type(nproc)} with ndim {nproc.ndim}")

    # read variables return a numpy array with corresponding selection
    steps = int(vars['physical_time']['AvailableStepsCount'])
    physical_time = s.read("physical_time", step_selection=[0, steps])
    print(
        f"physical_time is {physical_time} of type {type(physical_time)} with "
        f"ndim {physical_time.ndim} shape = {physical_time.shape}"
    )

    steps = int(vars['temperature']['AvailableStepsCount'])
    temperature = s.read("temperature", step_selection=[0, steps])
    temp_unit = s.read_attribute("temperature/unit")
    print(f"temperature array size is {temperature.size} of shape {temperature.shape}")
    print(f"temperature unit is {temp_unit} of type {type(temp_unit)}")

    steps = int(vars['pressure']['AvailableStepsCount'])
    pressure = s.read("pressure", step_selection=[0, steps])
    press_unit = s.read_attribute("pressure/unit")
    print()
```

```
$ python3 adios2-doc-read-filereader.py
variable_name: nproc      AvailableStepsCount: 1  Max: 4  Min: 4  Shape:      1
↳ SingleValue: true      Type: int64_t
variable_name: physical_time  AvailableStepsCount: 5  Max: 0.04      Min: 0  Shape:  1
↳ SingleValue: true      Type: double
variable_name: pressure      AvailableStepsCount: 5  Max: 1  Min: 1  Shape: 40
↳ SingleValue: false      Type: double
variable_name: temperature   AvailableStepsCount: 5  Max: 0  Min: 0  Shape: 40
↳ SingleValue: false      Type: double

nproc is 4 of type <class 'numpy.ndarray'> with ndim 0
physical_time is [0.  0.01 0.02 0.03 0.04] of type <class 'numpy.ndarray'> with ndim 1
↳ shape = (5,)
temperature array size is 200 of shape (200,)
temperature unit is K of type <class 'str'>
```



## 14.2 adios2 classes for Python

Stream is a high-level class that can perform most of the ADIOS functionality. FileReader is just a convenience class and is the same as Stream with “rra” (ReadRandomAccess) mode. FileReaders do not work with for loops as opposed to Streams that work step-by-step, rather one can access any step of any variable at will. The other classes, Adios, IO, Engine, Variable, Attribute and Operator correspond to the C++ classes. One needs to use them to extend the capabilities of the Stream class (e.g. using an external XML file for runtime configuration, changing the engine for the run, setting up a compression operator for an output variable, etc.)

**class** `adios2.Stream(path, mode, comm=None)`

High level implementation of the Stream class from the core API

### property `adios`

Adios instance associated to this Stream

**all\_blocks\_info**(*name*)

Extracts all available blocks information for a particular variable. This can be an expensive function, memory scales up with metadata: steps and blocks per step

### Args:

*name* (str): variable name

### Returns:

list of dictionaries with information of each step

**available\_attributes**(*varname=""*, *separator="/"*)

Returns a 2-level dictionary with attribute information. Read mode only.

### Parameters

#### **variable\_name**

If *varname* is set, attributes assigned to that variable are returned. The keys returned are attribute names with the prefix of *varname* + *separator* removed.

#### **separator**

concatenation string between *variable\_name* and attribute e.g. *varname* + *separator* + *name* (“var/attr”) Not used if *varname* is empty

### Returns

#### **attributes dictionary**

##### **key**

attribute name

##### **value**

attribute information dictionary

**available\_variables**()

Returns a 2-level dictionary with variable information. Read mode only.

### Parameters

#### **keys**

list of variable information keys to be extracted (case insensitive)  
`keys=['AvailableStepsCount','Type','Max','Min','SingleValue','Shape']` `keys=['Name']` returns only the variable names as 1st-level keys leave empty to return all possible keys

### Returns

#### **variables dictionary**

**key**

variable name

**value**

variable information dictionary

**begin\_step**(\*, *timeout=-1.0*)

Write mode: declare the starting of an output step. Pass data in `stream.write()` and `stream.write_attribute()`. All data will be published in `end_step()`.

Read mode: in streaming mode releases the current step (no effect in file based engines)

**close**()

Closes stream, thus becoming unreachable. Not required if using open in a with-as statement. Required in all other cases per-open to avoid resource leaks.

**current\_step**()

Inspect current step of the stream. The steps run from 0.

Note that in a real stream, steps from the producer may be missed if the consumer is slow and the producer is told to discard steps when no one is reading them in time. You may see non-consecutive numbers from this function call in this case.

Use `loop_index()` to get a loop counter in a `for ... .steps()` loop.

**Returns**

current step

**define\_variable**(*name*)

Define new variable without specifying its type and content. This only works for string output variables

**end\_step**()

Write mode: declaring the end of an output step. All data passed in `stream.write()` and all attributes passed in `stream.write_attribute()` will be published for consumers.

Read mode: in streaming mode releases the current step (no effect in file based engines)

**property engine**

Engine instance associated to this Stream

**inquire\_attribute**(*name*, *variable\_name=""*, *separator="/"*)

Inquire an attribute

**Parameters**

**name**

attribute name

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between `variable_name` and attribute e.g. `variable_name + separator + name` ("var/attr") Not used if `variable_name` is empty

**Returns**

The attribute if it is defined, otherwise None

**inquire\_variable**(*name*)

Inquire a variable

**Parameters**

**name**

variable name

**Returns**

The variable if it is defined, otherwise None

**property io**

IO instance associated to this Stream

**loop\_index()**

Inspect the loop counter when using for-in loops. This function returns consecutive numbers from 0.

**Returns**

the loop counter

**property mode**

Selected open mode

**num\_steps()**

READ MODE ONLY. Return the number of steps available. Note that this is the steps of a file/stream. Each variable has its own steps, which needs to be inspected with `var=stream.inquire_variable()` and then with `var.steps()`

**read**(*variable*: [Variable](#), *start*=[], *count*=[], *block\_id*=None, *step\_selection*=None)

**read**(*name*: str, *start*=[], *count*=[], *block\_id*=None, *step\_selection*=None)

Read a variable. Random access read allowed to select steps.

**Parameters****variable**

adios2.Variable object to be read Use `variable.set_selection()`, `set_block_selection()`, `set_step_selection()` to prepare a read

**start**

variable offset dimensions

**count**

variable local dimensions from offset

**block\_id**

(int) Required for reading local variables, local array, and local value.

**step\_selection**

(list): On the form of [start, count].

**Returns****array**

resulting array from selection

**read\_attribute**(*name*, *variable\_name*="", *separator*='/')

Reads a numpy based attribute

**Parameters****name**

attribute name

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between variable\_name and attribute e.g. variable\_name + separator + name (var/attr) Not used if variable\_name is empty

**Returns****array**

resulting array attribute data

**read\_attribute\_string**(name, variable\_name="", separator='/')

Reads a numpy based attribute

**Parameters****name**

attribute name

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between variable\_name and attribute e.g. variable\_name + separator + name (var/attr) Not used if variable\_name is empty

**Returns****array**

resulting array attribute data

**set\_parameters**(\*\*kwargs)

Sets parameters using a dictionary. Removes any previous parameter.

**Parameters****parameters**

input key/value parameters

**value**

parameter value

**step\_status**()

Inspect the stream status. Return adios2.bindings.StepStatus

**steps**(num\_steps=0, \*, timeout=-1.0)

Returns an iterator that can be use to iterate throught the steps. In each iteration begin\_step() and end\_step() will be internally called.

Write Mode: num\_steps is a mandatory argument and should specify the number of steps.

Read Mode: if num\_steps is not specified there will be as much iterations as provided by the actual engine. If num\_steps is given and there is not that many steps in a file/stream, an error will occur.

IMPORTANT NOTE: Do not use with ReadRandomAccess mode.

**write**(variable: [Variable](#), content)

**write**(name: str, content, shape=[], start=[], count=[], operations=None)

Writes a variable. Note that the content will be available for consumption only at the end of the for loop or in the end\_step() call.

**Parameters**

**variable**

adios2.Variable object to be written Use `variable.set_selection()`, `set_shape()`, `add_operation_string()` to prepare a write

**content**

variable data values

**write\_attribute**(*name*, *content*, *variable\_name*="", *separator*='/')

writes a self-describing single value array (numpy) variable

**Parameters****name**

attribute name

**array**

attribute numpy array data

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between *variable\_name* and attribute e.g. *variable\_name* + *separator* + *name* ("var/attr") Not used if *variable\_name* is empty

**class** `adios2.FileReader`(*path*, *comm*=None)

High level implementation of the FileReader class for read Random access mode

**class** `adios2.Adios`(*config\_file*=None, *comm*=None)

High level representation of the ADIOS class in the `adios2.bindings`

**at\_io**(*name*)

Inquire IO instance

**Args:**

*name* (str): IO instance name

**Returns:**

IO: IO instance

**declare\_io**(*name*)

Declare IO instance

**Args:**

*name* (str): IO instance name

**define\_operator**(*name*, *kind*, *parameters*={})

Add an operation (operator).

**Args:**

*name* (str): name of the operator. *kind* (str): name type of the operation. *params* (dict): parameters as a form of a dict for the operation.

**flush\_all**()

Flush all IO instances

**property impl**

Bindings implementation of the class

**inquire\_operator**(*name*)

Add an operation (operator).

**Args:**

name (str): name of the operator.

**Return:**

Operator: requested operator

**remove\_all\_ios**()

Remove all IO instance

**remove\_io**(*name*)

Remove IO instance

**Args:**

name (str): IO instance name

**class** adios2.**IO**(*impl, name, adiosobj*)

High level representation of the IO class in the adios2.bindings

**add\_transport**(*kind, parameters={}*)

Adds a transport and its parameters to current IO. Must be supported by current engine type.

**Parameters**

**kind**

must be a supported transport type for current engine.

**parameters**

acceptable parameters for a particular transport CAN'T use the keywords "Transport" or "transport" in key

**Returns**

**transport\_index**

handler to added transport

**adios**()

Adios instance associated to this IO

**available\_attributes**(*varname="", separator='/'*)

Returns a 2-level dictionary with attribute information. Read mode only.

**Parameters**

**variable\_name**

If varname is set, attributes assigned to that variable are returned. The keys returned are attribute names with the prefix of varname + separator removed.

**separator**

concatenation string between variable\_name and attribute e.g. varname + separator + name ("var/attr") Not used if varname is empty

**Returns**

**attributes dictionary**

**key**

attribute name

**value**

attribute information dictionary

**available\_variables()**

Returns a 2-level dictionary with variable information. Read mode only.

**Parameters****keys**

list of variable information keys to be extracted (case insensitive)  
keys=['AvailableStepsCount','Type','Max','Min','SingleValue','Shape'] keys=['Name'] returns only the variable names as 1st-level keys leave empty to return all possible keys

**Returns****variables dictionary****key**

variable name

**value**

variable information dictionary

**define\_attribute**(name, content=None, variable\_name="", separator='/')

Define an Attribute

**Parameters****name**

attribute name

**content**

attribute numpy array data

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between variable\_name and attribute e.g. variable\_name + separator + name ("var/attr") Not used if variable\_name is empty

**define\_variable**(name, content=None, shape=[], start=[], count=[], is\_constant\_dims=False)

writes a variable

**Parameters****name**

variable name

**content**

variable data values

**shape**

variable global MPI dimensions.

**start**

variable offset for current MPI rank.

**count**

variable dimension for current MPI rank.

**isConstantDims**

Whether dimensions are constant

**engine\_type()**

Return engine type

**flush\_all()**

Flush all engines attached to this IO instance

**property impl**

Bindings implementation of the class

**inquire\_attribute**(*name*, *variable\_name*="", *separator*='/')

Inquire an Attribute

**Parameters****name**

attribute name

**variable\_name**

if attribute is associated with a variable

**separator**

concatenation string between variable\_name and attribute e.g. variable\_name + separator + name ("var/attr") Not used if variable\_name is empty

**inquire\_variable**(*name*)

Inquire a variable

**Parameters****name**

variable name

**Returns**

The variable if it is defined, otherwise None

**open**(*name*, *mode*, *comm*=None)

Open an engine

**Parameters****name**

Engine name

**mode**

engine mode

**comm**

MPI communicator, optional

**parameters()**

Return parameter associated to this io instance

**Return:**

dict: str->str parameters

**remove\_all\_attributes()**

Remove all define attributes

**remove\_all\_variables()**

Remove all variables in the IO instance

**remove\_attribute**(*name*)

Remove an Attribute

**Parameters**



**name**  
attribute name

**remove\_variable**(*name*)  
Remove a variable

**Parameters**

**name**  
Variable name

**set\_engine**(*name*)  
Set engine for this IO instance

**Args:**  
name (str): name of engine

**set\_parameter**(*key*, *value*)  
Set a parameter for this IO instance

**Args:**  
key (str): value (str):

**set\_parameters**(*parameters*)  
Set parameters for this IO instance

**Args:**  
parameters (dict):

**class** adios2.**Engine**(*implementation*)  
High level representation of the Engine class in the adios2.bindings

**all\_blocks\_info**(*name*)  
Returns a list of BlocksInfo for a all steps

**Args:**  
name (str): variable name

**Returns:**  
list of BlockInfos

**begin\_step**(\**args*, \*\**kwargs*)  
Start step

**between\_step\_pairs**()  
End step

**blocks\_info**(*name*, *step*)  
Returns a BlocksInfo for a given step

**Args:**  
name (str): variable name step (int): step in question

**Returns:**  
list of dicts describing each BlockInfo

**close**(*transport\_index=-1*)  
Close an transports of the engine

**Args:**  
transportIndex (int): if -1 close all transports

**current\_step()**

Returns the current step

**end\_step()**

End step

**flush(*transport\_index=-1*)**

Flush all transports attached to this Engine instance

**get(*variable, content=None, mode=<Mode.Sync: 6>*)**

Gets the content of a variable

**Parameters****name**

variable name

**content**

output variable data values

**mode**

when to perform the communication, (Deferred or Asynchronous).

**Returns**

Content of the variable when the content argument is omitted.

**property impl**

Bindings implementation of the class

**lock\_reader\_selections()**

Locks the data selection for read

**lock\_writer\_definitions()**

Locks the data selection for write

**perform\_data\_write()**

Perform the transport data writes

**perform\_gets()**

Perform the gets calls

**perform\_puts()**

Perform the puts calls

**put(*variable, content, mode=<Mode.Deferred: 7>*)**

Puts content in a variable

**Parameters****name**

variable name

**content**

variable data values

**mode**

when to perform the communication, (Deferred or Asynchronous).

**steps()**

Returns number of available steps

**class** adios2.**Variable**(*implementation*)

High level representation of the Attribute class in the adios2.bindings

**add\_operation**(*operation*, *params*={})

Add an operation (operator).

**Args:**

name (Operator): name of the operation. params (dict): parameters as a form of a dict for the operation.

**add\_operation\_string**(*name*, *params*={})

Add an operation (operator) as a string

**Args:**

name (str): name of the operation. params (dict): parameters as a form of a dict for the operation.

**block\_id**()

BlockID of this variable.

**Returns:**

int: BlockID of this variable.

**count**()

Current selected count for this variable.

**Returns:**

int: Current selected count.

**property impl**

Bindings implementation of the class

**name**()

Name of the Variable

**Returns:**

str: Name of the Variable.

**operations**()

Current operations (operators) assigned to this Variable.

**Returns:**

list(Operators): operators assigned.

**remove\_operations**()

Remove operations (operators) assigned to this Variable.

**selection\_size**()

Current selection size selected for this variable.

**Returns:**

int: Current selection size selected.

**set\_block\_selection**(*block\_id*)

Set BlockID for this variable.

**Args:**

block\_id (int): Selected BlockID.

**set\_selection**(*selection*)

Set selection for this variable.

**Args:**

selection (list): list of the shape [[start], [count]], note that start and count can contain more than one element.

**set\_shape(*shape*)**

Set Shape (dimensions) for this variable

**Args:**

shape (list): desired shape (dimensions).

**set\_step\_selection(*step\_selection*)**

Set current step selection (For RRA or ReadRandomAccess)

**Args:**

step\_selection (list): On the form of [start, count].

**shape(*step=None*)**

Get the shape assigned to the given step for this variable.

**Args:**

step (int): Desired step. Only in ReadRandomAccess mode

**Returns:**

list: shape of the specified step in the form of [start, count].

**shape\_id()**

Get the ShapeID assigned to this variable.

**Returns:**

int: ShapeID assigned to this variable.

**single\_value()**

Check if this variable is a single value.

**Returns:**

bool: true if this variable is a single value.

**sizeof()**

Size in bytes of the contents of the variable.

**Returns:**

int: size in bytes of the contents.

**start()**

The current selected start of the variable.

**Returns:**

int: The current selected start of the variable.

**steps()**

The number of steps of the variable. This is always 1 in a stream. In ReadRandomAccess mode, this function returns the total number of steps available, which can be used when selecting steps for read.

**Returns:**

int: The available steps of the variable.

**steps\_start()**

The available start step, this is needed variables can start at any time step. This is for ReadRandomAccess.

**Returns:**

int: the starting step of for this Variable.

**type()**

Type of the Variable

**Returns:**

str: Type of the Variable.

**class** adios2.**Attribute**(*io, name, \*args, \*\*kwargs*)

High level representation of the Attribute class in the adios2.bindings

**data()**

Content of the Attribute

**Returns:**

Content of the Attribute as a non string.

**data\_string()**

Content of the Attribute

**Returns:**

Content of the Attribute as a str.

**property impl**

Bindings implementation of the class

**name()**

Name of the Attribute

**Returns:**

Name of the Attribute as a str.

**single\_value()**

True if the attribute is a single value, False if it is an array

**Returns:**

True or False.

**type()**

Type of the Attribute

**Returns:**

Type of the Attribute as a str.

**class** adios2.**Operator**(*implementation, name*)

High level representation of the Attribute class in the adios2.bindings

**get\_parameters()**

Get parameters associated to this Operator

**Returns:**

dict: parameters

**property impl**

Bindings implementation of the class

**set\_parameter**(*key, value*)

Set parameter associated to this Operator

**Args:**

str: key str: value

## 14.3 Python bindings to C++

**Note:** The bindings to the C++ functions is the basis of the native Python API described before. It is still accessible to users who used the “Full Python API” pre-2.10. In order to make old scripts working with 2.10 and later versions, change the import line in the python script.

---

```
import adios2.bindings as adios2
```

The full Python APIs follows very closely the full C++11 API interface. All of its functionality is now in the native API as well, so its use is discouraged for future scripts.

### 14.3.1 Examples using the Python bindings in the ADIOS2 repository

- **Simple file-based examples**
  - examples/hello/helloWorld/hello-world-bindings.py
  - examples/hello/bpReader/bpReaderHeatMap2D-bindings.py
  - examples/hello/bpWriter/bpWriter-bindings.py
- **Staging examples using staging engines SST and DataMan**
  - examples/hello/sstWriter/sstWriter-bindings.py
  - examples/hello/sstReader/sstReader-bindings.py

### 14.3.2 ADIOS class

**class** `adios2.bindings.ADIOS`

**AtIO**(*self*: `adios2.bindings.adios2_bindings.ADIOS`, *arg0*: *str*) → `adios2::py11::IO`

returns an IO object previously defined IO object with `DeclareIO`, throws an exception if not found

**DeclareIO**(*self*: `adios2.bindings.adios2_bindings.ADIOS`, *arg0*: *str*) → `adios2::py11::IO`

spawn IO object component returning a IO object with a unique name, throws an exception if IO with the same name is declared twice

**DefineOperator**(*self*: `adios2.bindings.adios2_bindings.ADIOS`, *arg0*: *str*, *arg1*: *str*, *arg2*: `Dict[str, str]`) → `adios2::py11::Operator`

**FlushAll**(*self*: `adios2.bindings.adios2_bindings.ADIOS`) → `None`

flushes all engines in all spawned IO objects

**InquireOperator**(*self*: `adios2.bindings.adios2_bindings.ADIOS`, *arg0*: *str*) → `adios2::py11::Operator`

**RemoveAllIOs**(*self*: `adios2.bindings.adios2_bindings.ADIOS`) → `None`

DANGER ZONE: remove all IOs in current ADIOS object, creates dangling objects to parameters, variable, attributes, engines created with removed IO

**RemoveIO**(*self*: `adios2.bindings.adios2_bindings.ADIOS`, *arg0*: *str*) → `bool`

DANGER ZONE: remove a particular IO by name, creates dangling objects to parameters, variable, attributes, engines created with removed IO

### 14.3.3 IO class

**class** adios2.bindings.IO

**AddTransport**(*self*: adios2.bindings.adios2\_bindings.IO, *type*: str, *parameters*: Dict[str, str] = {}) → int

**AvailableAttributes**(*self*: adios2.bindings.adios2\_bindings.IO, *varname*: str = "", *separator*: str = '/') → Dict[str, Dict[str, str]]

**AvailableVariables**(*self*: adios2.bindings.adios2\_bindings.IO) → Dict[str, Dict[str, str]]

**DefineAttribute**(\*args, \*\*kwargs)

Overloaded function.

1. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *array*: numpy.ndarray, *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
2. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *stringValue*: str, *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
3. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *strings*: List[str], *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
4. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *ints*: List[int], *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
5. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *doubles*: List[float], *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
6. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *complexes*: List[complex], *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute
7. DefineAttribute(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *value*: object, *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute

**DefineVariable**(\*args, \*\*kwargs)

Overloaded function.

1. DefineVariable(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *array*: numpy.ndarray, *shape*: List[int] = [], *start*: List[int] = [], *count*: List[int] = [], *isConstantDims*: bool = False) → adios2::py11::Variable
2. DefineVariable(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *value*: object, *shape*: List[int] = [], *start*: List[int] = [], *count*: List[int] = [], *isConstantDims*: bool = False) → adios2::py11::Variable
3. DefineVariable(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str) → adios2::py11::Variable

**EngineType**(*self*: adios2.bindings.adios2\_bindings.IO) → str

**FlushAll**(*self*: adios2.bindings.adios2\_bindings.IO) → None

**InquireAttribute**(*self*: adios2.bindings.adios2\_bindings.IO, *name*: str, *variable\_name*: str = "", *separator*: str = '/') → adios2::py11::Attribute

**InquireVariable**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str) → adios2::py11::Variable

**Open**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str, *arg1*: int) → adios2::py11::Engine

**Parameters**(*self*: adios2.bindings.adios2\_bindings.IO) → Dict[str, str]

**RemoveAllAttributes**(*self*: adios2.bindings.adios2\_bindings.IO) → None

**RemoveAllVariables**(*self*: adios2.bindings.adios2\_bindings.IO) → None

**RemoveAttribute**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str) → bool

**RemoveVariable**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str) → bool

**SetEngine**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str) → None

**SetParameter**(*self*: adios2.bindings.adios2\_bindings.IO, *arg0*: str, *arg1*: str) → None

**SetParameters**(*self*: adios2.bindings.adios2\_bindings.IO, *parameters*: Dict[str, str] = {}) → None

### 14.3.4 Variable class

**class** adios2.bindings.Variable

**AddOperation**(\*args, \*\*kwargs)

Overloaded function.

1. AddOperation(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: adios2.bindings.adios2\_bindings.Operator, *arg1*: Dict[str, str]) → int
2. AddOperation(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: str, *arg1*: Dict[str, str]) → int

**BlockID**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**Count**(*self*: adios2.bindings.adios2\_bindings.Variable) → List[int]

**Name**(*self*: adios2.bindings.adios2\_bindings.Variable) → str

**Operations**(*self*: adios2.bindings.adios2\_bindings.Variable) → List[adios2.bindings.adios2\_bindings.Operator]

**RemoveOperations**(*self*: adios2.bindings.adios2\_bindings.Variable) → None

**SelectionSize**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**SetBlockSelection**(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: int) → None

**SetSelection**(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: Tuple[List[int], List[int]]) → None

**SetShape**(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: List[int]) → None

**SetStepSelection**(*self*: adios2.bindings.adios2\_bindings.Variable, *arg0*: Tuple[int, int]) → None

**Shape**(*self*: adios2.bindings.adios2\_bindings.Variable, *step*: int = 18446744073709551615) → List[int]

**ShapeID**(*self*: adios2.bindings.adios2\_bindings.Variable) → adios2.bindings.adios2\_bindings.ShapeID

**SingleValue**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**Sizeof**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**Start**(*self*: adios2.bindings.adios2\_bindings.Variable) → List[int]

**Steps**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**StepsStart**(*self*: adios2.bindings.adios2\_bindings.Variable) → int

**Type**(*self*: adios2.bindings.adios2\_bindings.Variable) → str



### 14.3.5 Attribute class

**class** adios2.bindings.Attribute

**Data**(self: adios2.bindings.adios2\_bindings.Attribute) → numpy.ndarray

**DataString**(self: adios2.bindings.adios2\_bindings.Attribute) → List[str]

**Name**(self: adios2.bindings.adios2\_bindings.Attribute) → str

**SingleValue**(self: adios2.bindings.adios2\_bindings.Attribute) → bool

**Type**(self: adios2.bindings.adios2\_bindings.Attribute) → str

### 14.3.6 Engine class

**class** adios2.bindings.Engine

**BeginStep**(\*args, \*\*kwargs)

Overloaded function.

1. **BeginStep**(self: adios2.bindings.adios2\_bindings.Engine, mode: adios2.bindings.adios2\_bindings.StepMode, timeoutSeconds: float = -1.0) -> adios2.bindings.adios2\_bindings.StepStatus
2. **BeginStep**(self: adios2.bindings.adios2\_bindings.Engine) -> adios2.bindings.adios2\_bindings.StepStatus

**BetweenStepPairs**(self: adios2.bindings.adios2\_bindings.Engine) → bool

**BlocksInfo**(self: adios2.bindings.adios2\_bindings.Engine, arg0: str, arg1: int) → List[Dict[str, str]]

**Close**(self: adios2.bindings.adios2\_bindings.Engine, transportIndex: int = -1) → None

**CurrentStep**(self: adios2.bindings.adios2\_bindings.Engine) → int

**EndStep**(self: adios2.bindings.adios2\_bindings.Engine) → None

**Flush**(self: adios2.bindings.adios2\_bindings.Engine, arg0: int) → None

**Get**(\*args, \*\*kwargs)

Overloaded function.

1. **Get**(self: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, array: numpy.ndarray, launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Deferred: 7>) -> None
2. **Get**(self: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Deferred: 7>) -> str

**LockReaderSelections**(self: adios2.bindings.adios2\_bindings.Engine) → None

**LockWriterDefinitions**(self: adios2.bindings.adios2\_bindings.Engine) → None

**Name**(self: adios2.bindings.adios2\_bindings.Engine) → str

**PerformDataWrite**(self: adios2.bindings.adios2\_bindings.Engine) → None

**PerformGets**(self: adios2.bindings.adios2\_bindings.Engine) → None

**PerformPuts**(*self*: adios2.bindings.adios2\_bindings.Engine) → None

**Put**(\*args, \*\*kwargs)

Overloaded function.

1. Put(*self*: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, array: numpy.ndarray, launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Deferred: 7>) -> None
2. Put(*self*: adios2.bindings.adios2\_bindings.Engine, arg0: adios2.bindings.adios2\_bindings.Variable, arg1: str) -> None
3. Put(*self*: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, ints: List[int], launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Sync: 6>) -> None
4. Put(*self*: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, floats: List[float], launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Sync: 6>) -> None
5. Put(*self*: adios2.bindings.adios2\_bindings.Engine, variable: adios2.bindings.adios2\_bindings.Variable, complexes: List[complex], launch: adios2.bindings.adios2\_bindings.Mode = <Mode.Sync: 6>) -> None

**Steps**(*self*: adios2.bindings.adios2\_bindings.Engine) → int

**Type**(*self*: adios2.bindings.adios2\_bindings.Engine) → str

### 14.3.7 Operator class

**class** adios2.bindings.Operator

**Parameters**(*self*: adios2.bindings.adios2\_bindings.Operator) → Dict[str, str]

**SetParameter**(*self*: adios2.bindings.adios2\_bindings.Operator, arg0: str, arg1: str) → None

**Type**(*self*: adios2.bindings.adios2\_bindings.Operator) → str

### 14.3.8 Query class

**class** adios2.bindings.Query

**GetBlockIDs**(*self*: adios2.bindings.adios2\_bindings.Query) → List[int]

**GetResult**(*self*: adios2.bindings.adios2\_bindings.Query) → List[Tuple[List[int], List[int]]]

## 14.4 Transition from old API to new API

A python script using the high-level API of 2.9 and earlier needs to be modified to work with 2.10 and later.

- adios2.open() is replaced with adios2.Stream(), and does not have 4th and 5th optional arguments for external xml and IO name.
- the `for in file` is replaced with `for _ in file.steps()` but it works for both writing (by specifying the number of output steps) and reading (for the number of available steps in a stream/file).

```
# OLD API
import adios2

# NEW API
from adios2 import Adios, Stream

# NEW API: this still works
import adios2

# OLD API
fr = adios2.open(args.instream, "r", mpi.comm_app, "adios2.xml", "SimulationOutput")

# NEW API
adios = Adios("adios2.xml", mpi.comm_app)
io = adios.declare_io("SimulationOutput")
fr = Stream(io, args.instream, "r", mpi.comm_app)

# OLD API
for fr_step in fr:
    fr_step....

# NEW API 1
for _ in fr.steps():
    fr....

# NEW API 2
for fr_step in fr.steps():
    fr_step....
```



## AGGREGATION

The basic problem of large-scale I/O is that the N-to-1 and N-to-N (process-to-file) patterns do not scale and one must set the number of files in an output to the capability of the file system, not the size of the application. Hence,  $N$  processes need to write to  $M$  files to

- 1) utilize the bandwidth of the file system and to
- 2) minimize the cost of multiple process writing to a single file, while
- 3) not overwhelming the file system with too many files.

### 15.1 Aggregation in BP5

There are two implementations of aggregation in BP5, none of them is the same as the one in BP4. The aggregation setup in ADIOS2 consist of: a) *NumAggregators*, which processes do write to disk (others will send data to them), and b) *NumSubFiles*, how many files they will write.

**EveryoneWritesSerial** is a simple aggregation strategy. Every process is writing its own data to disk, to one particular file only, and the processes are serialized over each particular file. In this aggregator, *NumAggregators* = *NumSubFiles* (=  $M$ ). This approach should scale well with application size. On Summit's GPFS though we observe that a single writer per compute node is better than multiple process writing to the file system, hence this aggregation method performs poorly there.

**EveryoneWrites** is the same strategy as the previous except that every process immediately writes its own data to its designated file. Since it basically implements an N-to-N write pattern, this method does not scale, so only use it up to a moderate number of processes (1-4 process \* number of file system servers). At small scale, as long as the file system can deal with the on-rush of the write requests, this method can provide the fastest I/O.

**TwoLevelShm** has a subset of processes that actually write to disk (*NumAggregators*). There must be at least one process per compute node, which creates a shared-memory segment for other processes on the node to send their data. The aggregator process basically serializes the writing of data from this subset of processes (itself and the processes that send data to it). TwoLevelShm performs similarly to EveryoneWritesSerial on Lustre, and is the only good option on Summit's GPFS.

The number of files (*NumSubFiles*) can be smaller than *NumAggregators*, and then multiple aggregators will write to one file concurrently. Such a setup becomes useful when the number of nodes is many times more than the number of file servers.

TwoLevelShm works best if each process's output data fits into the shared-memory segment, which holds two pages. Since POSIX writes are limited to about 2GB, the best setup is to use 4GB shared-memory size by each aggregator. This is the default size, but you can use the *MaxShmSize* parameter to set this lower if necessary. At runtime, BP5 will only allocate twice the maximum size of the largest data size any process has, but up to *MaxShmSize*. If the data from two processes does not fit into the shared-memory segment, BP5 will need to perform multiple iterations of copy and disk-write, which is generally slower than writing large data blocks at once.

The **default setup** is *TwoLevelShm*, where *NumAggregators* is the number of compute nodes the application is running on, and the number of files is the same. This setup is good for Summit's GPFS and good for Lustre at large scale. However, the default setup leaves potential performance on the table when running applications at smaller scale, where the one process per node setup cannot utilize the full bandwidth of a large parallel file system.

## MEMORY MANAGEMENT

### 16.1 BP4 buffering

BP4 has a simple buffering mechanism to provide ultimate performance at the cost of high memory usage: all user data (passed in *Put()* calls) is buffered in one contiguous memory allocation and writing/aggregation is done with this large buffer in *EndStep()*. Aggregation in BP4 uses MPI to send this buffer to the aggregator and hence maintaining two such large buffers. Basically, if an application writes  $N$  bytes of data in a step, then BP4 needs approximately  $2xN$  bytes extra memory for buffering.

A potential performance problem is that BP4 needs to extend the buffer occasionally to fit more incoming data (more *Put()* calls). At large sizes the reallocation may need to move the buffer into a different place in memory, which requires copying the entire existing buffer. When there are GBs of data already buffered, this copy will noticeably decrease the overall observed write performance. This situation can be avoided if one can guess a usable upper limit to how much data each process is going to write, and telling this to the BP4 engine through the **InitialBufferSize** parameter before *Open()*.

Another potential problem is that reallocation may fail at some point, well before the limits of memory, since it needs a single contiguous allocation be available.

### 16.2 BP5 buffering

BP5 is designed to use less memory than BP4. The buffer it manages is a list of large chunks. The advantages of the list of chunks is that no reallocation of existing buffer is needed, and that BP5 can potentially allocate more buffer than BP4 since it requests many smaller chunks instead of a large contiguous buffer. In general, chunks should be as big as the system/application can afford, up to **2147381248** bytes (almost but less than 2GB, the actual size limit POSIX *write()* calls have). Each chunk will result in a separate write call, hence minimizing the number of chunks is preferred. The current default is set to 128MB, so please increase this on large computers if you can and if you write more than that amount of data per process, using the parameter **BufferChunkSize**.

Second, BP5 can add a large user variable as a chunk to this list without copying it at all and use it directly to write (or send to aggregator). *Put(..., adios2::Mode::Deferred)* will handle the user data directly, unless its size is below a threshold (see parameter **MinDeferredSize**).

---

**Note:** Do not call *PerformPuts()* when using BP5, because this call forces copying all user data into the internal buffer before writing, eliminating all benefits of zero-copy that BP5 provides when operating with large buffers. Instead, consider using *Put()* with the Sync option if you want to force ADIOS to copy data immediately. Alternatively, BP5 offers *PerformDataWrite()*, an collective operation that actually moves data to storage, potentially freeing up buffer and application memory.

---

Third, BP5 is using a shared memory segment on each compute node for aggregation, instead of MPI. The best settings for the shared memory is 4GB (see parameter **MaxShmSize**), enough place for two chunks with the POSIX write limit. More is useless but can be smaller if a system/application cannot allow this much space for aggregation (but there will be more write calls to disk as a result).

## 16.3 Span object in internal buffer

Another option to decrease memory consumption is to pre-allocate space in the BP4/BP5 buffer and then prepare output variables directly in that space. This will avoid a copy and the need for doubling memory for temporary variables that are only created for output purposes. This Span feature is only available in C++. See the *Span()* function in [\*Engine class\*](#) `./api_full/api_full.html#engine-class`



## GPU-AWARE I/O

The Put and Get functions in the default file engine (BP5) and some streaming engines (SST, DataMan) can receive user buffers allocated on the host or the device in both Sync and Deferred modes.

---

**Note:** Buffers allocated on the device with CUDA, HIP and SYCL are supported.

---

If ADIOS2 is built without GPU support, only buffers allocated on the host are supported. When GPU support is enabled, the default behavior is for ADIOS2 to automatically detect where the buffer memory physically resides.

Users can also provide information about where the buffer was allocated by using the SetMemorySpace function within each variable.

```
enum class MemorySpace
{
    Detect, ///Detect the memory space automatically
    Host,   ///Host memory space (default)
    GPU     ///GPU memory spaces
};
```

If ADIOS2 is built without GPU support, the available MemorySpace values are only Detect and Host.

ADIOS2 can use a CUDA or Kokkos backend for enabling GPU support. Only one backend can be active at a given time based on how ADIOS2 is build.

### 17.1 Building ADIOS2 with a GPU backend

By default both backends are OFF even if CUDA or Kokkos are installed and available to avoid a possible conflict between if both backends are enabled at the same time.

#### 17.1.1 Building with CUDA enabled

The ADIOS2 default behavior is to turn OFF the CUDA backend. Building with the CUDA backend requires -DADIOS2\_USE\_Kokkos=ON and an available CUDA toolkit on the system.

When building ADIOS2 with CUDA enabled, the user is responsible with setting the correct CMAKE\_CUDA\_ARCHITECTURES (e.g. for Summit the CMAKE\_CUDA\_ARCHITECTURES needs to be set to 70 to match the NVIDIA Volta V100).

### 17.1.2 Building with Kokkos enabled

The Kokkos library can be used to enable GPU within ADIOS2. Based on how Kokkos is build, either the CUDA, HIP or SYCL backend will be enabled. Building with Kokkos requires `-DADIOS2_USE_Kokkos=ON`. The `CMAKE_CUDA_ARCHITECTURES` is set automatically to point to the same architecture used when configuring the Kokkos library.

---

**Note:** Kokkos version  $\geq 3.7$  is required to enable the GPU backend in ADIOS2

---

## 17.2 Writing GPU buffers

The ADIOS2 API for Device pointers is identical to using Host buffers for both the read and write logic. Internally each ADIOS2 variable holds a memory space for the data it receives. Once the memory space is set (either directly by the user through calls to `SetMemorySpace` or after detecting the buffer memory space the first Put or Get call) to either Host or Device, it cannot be changed.

The `examples/hello` folder contains several codes that use Device buffers:

- `bpStepsWriteRead{Cuda|Hip}` show CUDA and HIP codes using BP5 with GPU pointers
- `bpStepsWriteReadKokkos` contains Fortran and C++ codes using `Kokkos::View` with different memory spaces and a Kokkos code using different layouts on Host buffers
- `datamanKokkos` shows an example of streaming a `Kokkos::View` with DataMan using different memory spaces
- `sstKokkos` shows an example of streaming a `Kokkos::View` with SST using different memory spaces

### 17.2.1 Example using a Device buffer

The following is a simple example of writing data to storage directly from a GPU buffer allocated with CUDA relying on the automatic detection of device pointers in ADIOS2.

```
float *gpuSimData;
cudaMalloc(&gpuSimData, N * sizeof(float));
cudaMemset(gpuSimData, 0, N);
auto data = io.DefineVariable<float>("data", shape, start, count);

io.SetEngine("BP5");
adios2::Engine bpWriter = io.Open(fname, adios2::Mode::Write);
// Simulation steps
for (size_t step = 0; step < nSteps; ++step)
{
    bpWriter.BeginStep();
    bpWriter.Put(data, gpuSimData, adios2::Mode::Deferred); // or Sync
    bpWriter.EndStep();
}
```

If the `SetMemorySpace` function is used, the ADIOS2 library will not detect automatically where the buffer was allocated and will use the information provided by the user for all subsequent Puts or Gets. Example:

```
data.SetMemorySpace(adios2::MemorySpace::GPU);
for (size_t step = 0; step < nSteps; ++step)
{
    bpWriter.BeginStep();
    bpWriter.Put(data, gpuSimData, adios2::Mode::Deferred); // or Sync
    bpWriter.EndStep();
}
```

Underneath, ADIOS2 relies on the backend used at build time to transfer the data. If ADIOS2 was build with CUDA, only CUDA buffers can be provided. If ADIOS2 was build with Kokkos (with CUDA enabled) only CUDA buffers can be provided. If ADIOS2 was build with Kokkos (with HIP enabled) only HIP buffers can be provided.

**Note:** The SYCL backend in Kokkos can be used to run on Nvidia, AMD and Intel GPUs, but we recommend using SYCL for Intel, HIP for AMD and CUDA for Nvidia.

### 17.2.2 Kokkos applications

ADIOS2 supports GPU buffers provided in the form of `Kokkos::View` directly in the Put/Get calls. The memory space is automatically detected from the View information. In addition to the memory space, for `Kokkos::View` ADIOS2 also extracts the layout of the array and adjust the variable dimensions to be able to build the global shape (across ranks) of the array.

```
Kokkos::View<float *, Kokkos::CudaSpace> gpuSimData("data", N);
bpWriter.Put(data, gpuSimData);
```

If the CUDA backend is being used (and not Kokkos) to enable GPU support in ADIOS2, Kokkos applications can still directly pass `Kokkos::View` as long as the correct external header is included: `#include <adios2/cxx11/KokkosView.h>`.

## 17.3 Reading GPU buffers

The GPU-aware backend allows different layouts for global arrays without requiring the user to update the code for each case. The user defines the shape of the global array and ADIOS2 adjusts the dimensions for each rank according to the buffer layout and memory space.

The following example shows a global array of shape (4, 3) when running with 2 ranks, each contributing half of it.

```
Write on LayoutRight, read on LayoutRight
1 1 1 // rank 0
2 2 2
3 3 3 // rank 1
4 4 4
Write on LayoutRight, read on LayoutLeft
1 2 3 4
1 2 3 4
1 2 3 4
```

On the read side, the Shape function can take a memory space or a layout to return the correct dimensions of the variable. For the previous example, if a C++ code using two ranks wants to read the data into a GPU buffer, the Shape

of the local array should be (3, 2). If the same data will be read on CPU buffers, the shape should be (2, 3). Both of the following code would give acceptable answers:

```
auto dims_host = data.Shape(adios2::MemorySpace::Host);  
auto dims_device = data.Shape(adios2::ArrayOrdering::ColumnMajor);
```

## 17.4 Build scripts

The *scripts/build\_scripts* folder contains scripts for building ADIOS2 with CUDA or Kokkos backends for several DOE system: Summit (OLCF Nvidia), Crusher (OLCFi AMD), Perlmutter (NERSC Nvidia), Polaris (ALCF Nvidia).

---

**Note:** Perlmutter requires Kokkos  $\geq 4.0$

---

## ADIOS2 QUERY API

The query API in ADIOS2 allows a client to pass a query in XML or json format, and get back a list of blocks or sub-blocks that contains hits. Both BP4 and BP5 engines are supported.

### 18.1 The interface

User is expected to pass a query file (configFile), and init a read engine (engine) to construct a query and evaluate using the engine. (note that the engine and query should be using the same ADIOS IO)

```
class QueryWorker
{
public:
    // configFile has query, can be either xml or json
    QueryWorker(const std::string &configFile, adios2::Engine &engine);

    // touched_blocks is a list of regions specified by (start, count),
    // that contains data that satisfies the query file
    void GetResultCoverage(std::vector<adios2::Box<adios2::Dims>> &touched_blocks);
    ...
}
```

#### 18.1.1 A Sample Compound Query

This query targets a 1D variable “doubleV”, data of interest is  $(x > 6.6)$  or  $(x < -0.17)$  or  $(2.8 < x < 2.9)$  In addition, this query also specied an output region [start=5,count=80].

```
<adios-query>
  <io name="query">
  <var name="doubleV">
    <boundingbox start="5" count="80"/>
    <op value="OR">
      <range compare="GT" value="6.6"/>
      <range compare="LT" value="-0.17"/>
      <op value="AND">
        <range compare="LT" value="2.9"/>
        <range compare="GT" value="2.8"/>
      </op>
    </op>
  </var>
```

(continues on next page)

(continued from previous page)

```
</io>
</adios-query>
```

## 18.2 Code EXAMPLES:

### 18.2.1 C++:

```
while (reader.BeginStep() == adios2::StepStatus::OK)
{
    adios2::QueryWorker w = adios2::QueryWorker(queryFile, reader);
    w.GetResultCoverage(touched_blocks);

    std::cout << " ... now can read out touched blocks ... size=" << touched_blocks.
    ↪size()
           << std::endl;
}
```

The Full C++ example is here:

<https://github.com/ornladios/ADIOS2/blob/master/examples/query/test.cpp>

### 18.2.2 Python:

```
while (reader.BeginStep() == adios2.StepStatus.OK):
# say only rank 0 wants to process result
var = [queryIO.InquireVariable("T")]

if (rank == 0):
    touched_blocks = w.GetResult()
    doAnalysis(reader, touched_blocks, var)
```

Full python example is here:

<https://github.com/ornladios/ADIOS2/blob/master/testing/adios2/bindings/python/TestQuery.py>

This example generates data, the query file (in xml) and runs the query, all in python.

## PLUGINS

ADIOS now has the ability for users to load their own engines and operators through the plugin interface. The basic steps for doing this are:

1. Write your plugin class, which needs to inherit from the appropriate `Plugin*Interface` class.
2. Build as a shared library and add the path to your shared library to the `ADIOS2_PLUGIN_PATH` environment variable.
3. Start using your plugin in your application.

These steps are discussed in further detail below.

## 19.1 Writing Your Plugin Class

### 19.1.1 Engine Plugin

Your engine plugin class needs to inherit from the `PluginEngineInterface` class in the `adios2/engine/plugin/PluginEngineInterface.h` header. Depending on the type of engine you want to implement, you'll need to override a number of methods that are inherited from the `adios2::core::Engine` class. These are briefly described in the following table. More detailed documentation can be found in `adios2/core/Engine.h`.

Method	Engine Type	Description
<code>BeginStep()</code>	Read/Write	Indicates the beginning of a step
<code>EndStep()</code>	Read/Write	Indicates the end of a step
<code>CurrentStep()</code>	Read/Write	Returns current step info
<code>DoClose()</code>	Read/Write	Close a particular transport
<code>Init()</code>	Read/Write	Engine initialization
<code>InitParameters()</code>	Read/Write	Initialize parameters
<code>InitTransports()</code>	Read/Write	Initialize transports
<code>PerformPuts()</code>	Write	Execute all deferred mode Put
<code>Flush()</code>	Write	Flushes data and metadata to a transport
<code>DoPut()</code>	Write	Implementation for Put
<code>DoPutSync()</code>	Write	Implementation for Put (Sync mode)
<code>DoPutDeferred()</code>	Write	Implementation for Put (Deferred Mode)
<code>PerformGets()</code>	Read	Execute all deferred mode Get
<code>DoGetSync()</code>	Read	Implementation for Get (Sync mode)
<code>DoGetDeferred()</code>	Read	Implementation for Get (Deferred Mode)

Examples showing how to implement an engine plugin can be found in `examples/plugins/engine`. An example write engine is `ExampleWritePlugin.h`, while an example read engine is in `ExampleReadPlugin.h`. The writer is a

simple file writing engine that creates a directory (called `ExamplePlugin` by default) and writes variable information to `vars.txt` and actual data to `data.txt`. The reader example reads the files output by the writer example.

In addition to implementing the methods above, you'll need to implement `EngineCreate()` and `EngineDestroy()` functions so ADIOS can create/destroy the engine object. Because of C++ name mangling, you'll need to use `extern "C"`. Looking at `ExampleWritePlugin.h`, this looks like:

```
extern "C" {

adios2::plugin::ExampleWritePlugin *
EngineCreate(adios2::core::IO &io, const std::string &name,
             const adios2::Mode mode, adios2::helper::Comm comm)
{
    return new adios2::plugin::ExampleWritePlugin(io, name, mode,
                                                  comm.Duplicate());
}

void EngineDestroy(adios2::plugin::ExampleWritePlugin * obj)
{
    delete obj;
}

}
```

### 19.1.2 Operator Plugin

Your operator plugin class needs to inherit from the `PluginOperatorInterface` class in the `adios2/operator/plugin/PluginOperatorInterface.h` header. There's three methods that you'll need to override from the `adios2::core::Operator` class, which are described below.

Method	Description
<code>Operate()</code>	Performs the operation, e.g., compress data
<code>InverseOperate()</code>	Performs the inverse operation, e.g., decompress data
<code>IsDataTypeValid()</code>	Checks that a given data type can be processed

An example showing how to implement an operator plugin can be found at `plugins/EncryptionOperator.h` and `plugins/EncryptionOperator.cpp`. This operator uses `libsodium` for encrypting and decrypting data.

In addition to implementing the methods above, you'll need to implement `OperatorCreate()` and `OperatorDestroy()` functions so ADIOS can create/destroy the operator object. Because of C++ name mangling, you'll need to use `extern "C"`. Looking at `EncryptionOperator`, this looks like:

```
extern "C" {

adios2::plugin::EncryptionOperator *
OperatorCreate(const adios2::Params &parameters)
{
    return new adios2::plugin::EncryptionOperator(parameters);
}

void OperatorDestroy(adios2::plugin::EncryptionOperator * obj)
{
}
```

(continues on next page)



(continued from previous page)

```

    delete obj;
}

}

```

## 19.2 Build Shared Library

To build your plugin, your CMake should look something like the following (using the plugin engine example described above):

```

find_package(ADIOS2 REQUIRED)
set(BUILD_SHARED_LIBS ON)
add_library(PluginEngineWrite
    ExampleWritePlugin.cpp
)
target_link_libraries(PluginEngineWrite adios2::cxx11 adios2::core)

```

When using the Plugin Engine, ADIOS will check for your plugin at the path specified in the `ADIOS2_PLUGIN_PATH` environment variable. If `ADIOS2_PLUGIN_PATH` is not set, and a path is not specified when loading your plugin (see below steps for using a plugin in your application), then the usual `dlopen` search is performed (see the [dlopen man page](#)).

---

**Note:** The `ADIOS2_PLUGIN_PATH` environment variable can contain multiple paths, which must be separated with a `..`.

---

When building on Windows, you will likely need to explicitly export the Create and Destroy symbols for your plugin, as symbols are invisible by default on Windows. To do this in a portable way across platforms, you can add something similar to the following lines to your `CMakeLists.txt`:

```

include(GenerateExportHeader)
generate_export_header(PluginEngineWrite BASE_NAME plugin_engine_write)
target_include_directories(PluginEngineWrite PUBLIC
    $<BUILD_INTERFACE:${CMAKE_CURRENT_BINARY_DIR}>
    $<INSTALL_INTERFACE:include>)

```

Then in your plugin header, you'll need to `#include "plugin_engine_write_export.h"`. Then edit your function definitions as follows:

```

extern "C" {

PLUGIN_ENGINE_WRITE_EXPORT adios2::plugin::ExampleWritePlugin *
    EngineCreate(adios2::core::IO &io, const std::string &name,
        const adios2::Mode mode, adios2::helper::Comm comm);

PLUGIN_ENGINE_WRITE_EXPORT void
    EngineDestroy(adios2::plugin::ExampleWritePlugin * obj);

}

```

## 19.3 Using Your Plugin in an Application

For both types of plugins, loading the plugin is done by setting the `PluginName` and `PluginLibrary` parameters in an `adios2::Params` object or `<parameter>` XML tag.

### 19.3.1 Engine Plugins

For engine plugins, this looks like:

```
adios2::ADIOS adios;
adios2::IO io = adios.DeclareIO("writer");
io.SetEngine("Plugin");
adios2::Params params;
params["PluginName"] = "WritePlugin";
params["PluginLibrary"] = "PluginEngineWrite";
// If the engine plugin has any other parameters, these can be added to
// the same params object and they will be forwarded to the engine
io.SetParameters(params);
```

Where “WritePlugin” is the name that ADIOS will use to keep track of the plugin, and “PluginEngineWrite” is the shared library name. At this point you can open the engine and use it as you would any other ADIOS engine. You also shouldn’t need to make any changes to your CMake files for your application.

The second option is using an ADIOS XML config file. If you’d like to load your plugins through an XML config file, the following shows an example XML when using Engine Plugins:

```
<adios-config>
  <io name="writer">
    <engine type="Plugin">
      <parameter key="PluginName" value="WritePlugin" />
      <parameter key="PluginLibrary" value="PluginEngineWrite" />
      <!-- any parameters needed for your plugin can be added here in the_
↪parameter tag -->
    </engine>
  </io>
  <io name="reader">
    <engine type="Plugin">
      <parameter key="PluginName" value="ReadPlugin" />
      <parameter key="PluginLibrary" value="PluginEngineRead" />
      <!-- any parameters needed for your plugin can be added here in the_
↪parameter tag -->
    </engine>
  </io>
</adios-config>
```

The examples `examples/plugins/engine/examplePluginEngine_write.cpp` and `examples/plugins/engine/examplePluginEngine_read.cpp` are an example of how to use the engine plugins described above.

### 19.3.2 Operator Plugins

For operator plugins, the code to use your plugin looks like:

```
// for an adios2::Variable<T> var
adios2::Params params;
params["PluginName"] = "MyOperator";
params["PluginLibrary"] = "EncryptionOperator";
// example param required for the EncryptionOperator
params["SecretKeyFile"] = "test-key";
var.AddOperation("plugin", params);
```

If you'd like to load your operator plugin through an XML config file, the following shows an example:

```
<adios-config>
  <io name="writer">
    <variable name="data">
      <operation type="plugin">
        <parameter key="PluginName" value="OperatorPlugin"/>
        <parameter key="PluginLibrary" value="EncryptionOperator" />
        <parameter key="SecretKeyFile" value="test-key" />
      </operation>
    </variable>
    <engine type="BP5">
      </engine>
    </io>
  </adios-config>
```

The examples `examples/plugins/operator/examplePluginOperator_write.cpp` and `examples/plugins/engine/examplePluginOperator_read.cpp` show an example of how to use the `EncryptionOperator` plugin described above.

**Note:** You don't need to add the `lib` prefix or the shared library ending (e.g., `.so`, `.dll`, etc.) when setting `PluginLibrary`. ADIOS will add these when searching for your plugin library. If you do add the prefix/suffix, ADIOS will still be able to find your plugin. It's also possible to put the full path to the shared library here, instead of using `ADIOS2_PLUGIN_PATH`.



## CAMPAIGN MANAGEMENT

The campaign management in ADIOS2 is for collecting basic information and metadata about a collection of ADIOS2 output files, from a single application run or multiple runs. The campaign archive is a single file (.ACA) that can be transferred to other locations. The campaign file can be opened by ADIOS2 and all the metadata can be processed (including the values of GlobalValue and LocalValue variables, or min/max of each Arrays at each step and decomposition/min/max of each block in an Array at each step). However, Get() operations will only succeed to read actual data of the arrays, if the data belonging to the campaign is either local or some mechanism for remote data access to the location of the data is set up in advance.

**Warning:** In 2.10, Campaign Management is just a first prototype and is included only for evaluation purposes. It will change substantially in the future and campaign files produced by this version will unlikely to be supported going forward.

### 20.1 The idea

Applications produce one or more output files in a single run. Subsequent analysis and visualization runs produce more output files. Campaign is a data organization concept one step higher than a file. A campaign archive includes information about multiple files, including the scalar variable's values and the min/max of arrays and the location of the data files (host and directory information). A science project can agree on how to organize their campaigns, i.e., how to name them, what files to include in a single campaign archive, how to distribute them, how to name the hosts where the actual data resides.

#### 20.1.1 Example

The Gray-Scott example, that is included with ADIOS2, in *examples/simulation/gray-scott*, has two programs, Gray-Scott and PDF-Calc. The first one produces the main output *gs.bp* which includes the main 3D variables *U* and *V*, and a checkpoint file *ckpt.bp* with a single step in it. PDF-Calc processes the main output and produces histograms on 2D slices of *U* and *V* (*U/bins* and *U/pdf*) in *pdf.bp*. A campaign can include all the three output files as they logically belong together.

```
# run application as usual
$ mpirun -n 4 adios2_simulations_gray-scott settings-files.json
$ ls -d *.bp
ckpt.bp gs.bp

$ adios2_campaign_manager.py create demoproject/frontier_gray-scott_100
```

(continues on next page)

(continued from previous page)

```

$ mpirun -n 3 adios2_simulations_gray-scott_pdf-calc gs.bp pdf.bp 1000
$ ls -d *.bp
ckpt.bp gs.bp pdf.bp

$ adios2_campaign_manager.py update demoproject/frontier_gray-scott_100

$ adios2_campaign_manager.py info demoproject/frontier_gray-scott_100
info archive
ADIOS Campaign Archive, version 1.0, created on 2024-04-01 10:44:11.644942
hostname = OLCF    longhostname = frontier.olcf.ornl.gov
    dir = /lustre/orion/csc143/proj-shared/demo/gray-scott
        dataset = ckpt.bp    created on 2024-04-01 10:38:19
        dataset = gs.bp      created on 2024-04-01 10:38:17
        dataset = pdf.bp     created on 2024-04-01 10:38:08

# The campaign archive is small compared to the data it points to
$ du -sh *.bp
7.9M    ckpt.bp
40M     gs.bp
9.9M    pdf.bp

$ du -sh /lustre/orion/csc143/proj-shared/adios-campaign-store/demoproject/frontier_gray-
↪scott_100.aca
97K     /lustre/orion/csc143/proj-shared/adios-campaign-store/demoproject/frontier_gray-
↪scott_100.aca

# ADIOS can list the content of the campaign archive
$ bpls -l demoproject/frontier_gray-scott_100
double   ckpt.bp/U      {4, 34, 34, 66} = 0.171103 / 1
double   ckpt.bp/V      {4, 34, 34, 66} = 1.71085e-19 / 0.438921
int32_t   ckpt.bp/step   scalar = 700
double   gs.bp/U        10*{64, 64, 64} = 0.090778 / 1
double   gs.bp/V        10*{64, 64, 64} = 8.24719e-63 / 0.515145
int32_t   gs.bp/step     10*scalar = 100 / 1000
double   pdf.bp/U/bins   10*{1000} = 0.0908158 / 0.999938
double   pdf.bp/U/pdf    10*{64, 1000} = 0 / 4096
double   pdf.bp/V/bins   10*{1000} = 8.24719e-63 / 0.514267
double   pdf.bp/V/pdf    10*{64, 1000} = 0 / 4096
int32_t   pdf.bp/step     10*scalar = 100 / 1000

# scalar over steps is available in metadata
$ bpls -l demoproject/frontier_gray-scott_100 -d pdf.bp/step -n 10
int32_t   pdf.bp/step     10*scalar = 100 / 1000
(0)       100 200 300 400 500 600 700 800 900 1000

# Array decomposition including min/max are available in metadata
$ bpls -l demoproject/frontier_gray-scott_100 -D gs.bp/V
double   gs.bp/V          10*{64, 64, 64} = 8.24719e-63 / 0.515145
step 0:
  block 0: [ 0:63, 0:31, 0:31] = 8.24719e-63 / 0.410653
  block 1: [ 0:63, 32:63, 0:31] = 8.24719e-63 / 0.410652
  block 2: [ 0:63, 0:31, 32:63] = 8.24719e-63 / 0.410653

```

(continues on next page)

(continued from previous page)

```

    block 3: [ 0:63, 32:63, 32:63] = 8.24719e-63 / 0.410653
    ...
    step 9:
    block 0: [ 0:63, 0:31, 0:31] = 3.99908e-09 / 0.441847
    block 1: [ 0:63, 32:63, 0:31] = 3.99931e-09 / 0.44192
    block 2: [ 0:63, 0:31, 32:63] = 3.99928e-09 / 0.441813
    block 3: [ 0:63, 32:63, 32:63] = 3.99899e-09 / 0.441796

# Array data is only available if data is local
$ ./bin/bpls -l demoproject/frontier_gray-scott_100 -d pdf.bp/U/bins
double pdf.bp/U/bins 10*{1000} = 0.0908158 / 0.999938
(0, 0) 0.93792 0.937982 0.938044 0.938106 0.938168 0.93823 0.938292 0.938354 0.
↪ 938416 0.938479
...
(9,990) 0.990306 0.991157 0.992007 0.992858 0.993708 0.994559 0.995409 0.99626 0.
↪ 99711 0.997961

```

## 20.2 Setup

There are three paths/names important in the campaign setup.

- *hostname* is the name detected by the `adios2_campaign_manager` when creating a campaign archive, however, it is better to define a specific name the project agrees upon (e.g. OLCF, NERSC, ALCF) that identifies the generic location of the data and then use that name later to specify the modes of remote data access (not available in this release).
- *campaignstorepath* is the directory where all the campaign archives are stored. This should be shared between project members in a center, and a private one on every member's laptop. It is up to the project to determine what file sharing / synchronization mechanism to use to sync this directories. [Rclone is a great command-line tool](#) to sync the campaign store with many cloud-based file sharing services and cloud instances.
- *cachepath* is the directory where ADIOS can unpack metadata from the campaign archive so that ADIOS engines can read them as if they were entirely local datasets. The cache only contains the metadata for now but in the future data that have already been retrieved by previous read requests will be stored here as well.

Use `~/config/adios2/adios2.yaml` to specify these options.

```

$ cat ~/.config/adios2/adios2.yaml

Campaign:
  active: true
  hostname: OLCF
  campaignstorepath: /lustre/orion/csc143/proj-shared/adios-campaign-store
  cachepath: /lustre/orion/csc143/proj-shared/campaign-cache
  verbose: 0

$ ls -R ~/dropbox/adios-campaign-store
/lustre/orion/csc143/proj-shared/adios-campaign-store/demoproject:
frontier_gray-scott_100.aca

$ adios2_campaign_manager.py list
demoproject/frontier_gray-scott_100.aca

```

## 20.3 Remote access

For now, we have one way to access data, through SSH port forwarding and running a remote server program to read in data on the remote host and to send back the data to the local ADIOS program. *adios2\_remote\_server* is included in the adios installation. You need to use the one built on the host.

Launch the server by SSH-ing to the remote machine, and specifying the 26200 port for forwarding. For example:

```
$ ssh -L 26200:dtm.olcf.ornl.gov:26200 -l <username> dtm.olcf.ornl.gov "<path_to_adios_
↪install>/bin/adios2_remote_server -v "
```

Assuming the campaign archive was synced to a local machine's campaign store under *csc143/demoproject*, now we can retrieve data:

```
$ adios2_campaign_manager.py list
csc143/demoproject/frontier_gray-scott_100.aca

$ bpls -l csc143/demoproject/frontier_gray-scott_100
double   ckpt.bp/U      {4, 34, 34, 66} = 0.171103 / 1
...
double   pdf.bp/U/bins  10*{1000} = 0.0908158 / 0.999938

# metadata is extracted to the local cachepath
$ du -sh /tmp/campaign/OLCF/csc143/demoproject/frontier_gray-scott_100.aca/*
20K      /tmp/campaign/OLCF/csc143/demoproject/frontier_gray-scott_100.aca/ckpt.bp
40K      /tmp/campaign/OLCF/csc143/demoproject/frontier_gray-scott_100.aca/gs.bp
32K      /tmp/campaign/OLCF/csc143/demoproject/frontier_gray-scott_100.aca/pdf.bp

# data is requested from the remote server
# read 16 values (4x4x4) from U from last step, from offset 30,30,30
$ bpls -l csc143/demoproject/frontier_gray-scott_100 -d gs.bp/U -s "-1,30,30,30" -c "1,
↪4,4,4" -n 4
double   gs.bp/U      10*{64, 64, 64}
slice (9:9, 30:33, 30:33, 30:33)
(9,30,30,30)  0.89189 0.899854 0.899854 0.891891
(9,30,31,30)  0.899851 0.908278 0.908278 0.899852
(9,30,32,30)  0.899849 0.908276 0.908277 0.899851
(9,30,33,30)  0.891885 0.899848 0.899849 0.891886
(9,31,30,30)  0.89985 0.908276 0.908276 0.899849
(9,31,31,30)  0.908274 0.916977 0.916977 0.908274
(9,31,32,30)  0.908273 0.916976 0.916976 0.908273
(9,31,33,30)  0.899844 0.908271 0.908271 0.899844
(9,32,30,30)  0.89985 0.908276 0.908275 0.899848
(9,32,31,30)  0.908274 0.916976 0.916976 0.908272
(9,32,32,30)  0.908272 0.916975 0.916974 0.908271
(9,32,33,30)  0.899844 0.90827 0.90827 0.899842
(9,33,30,30)  0.89189 0.899851 0.899851 0.891886
(9,33,31,30)  0.89985 0.908275 0.908275 0.899847
(9,33,32,30)  0.899848 0.908274 0.908273 0.899845
(9,33,33,30)  0.891882 0.899845 0.899844 0.89188
```



## 20.4 Requirements

The Campaign Manager uses SQLite3 and ZLIB for its operations, and Python3 3.8 or higher for the *adios2\_campaign\_manager* tool. Check *bpls -Vv* to see if *CAMPAIGN* is in the list of “Available features”.

## 20.5 Limitations

- The Campaign Reader engine only supports ReadRandomAccess mode, not step-by-step reading. Campaign management will need to change in the future to support sorting the steps from different outputs to a coherent order.
- Updates to moving data for other location is not supported yet



## ADIOS2 IN ECP HARDWARE

ADIOS2 is widely used in ECP (Exascale Computing Project) HPC (high performance computing) systems, some particular ADIOS2 features needs from specifics workarounds to run successfully.

### 21.1 OLCF CRUSHER

#### 21.1.1 SST MPI Data Transport

MPI Data Transport relies on client-server features of MPI which are currently supported in Cray-MPI implementations with some caveats. Here are some of the observed issues and what its workaround (if any) are:

**MPI\_Finalize** will block the system process in the “Writer/Producer” ADIOS2 instance. The reason is that the Producer ADIOS instance internally calls *MPI\_Open\_port* which somehow even after calling *MPI\_Close\_port* *MPI\_Finalize* still consider its port to be in used, hence blocking the process. The workaround is to use a *MPI\_Barrier(MPI\_COMM\_WORLD)* instead of *MPI\_Finalize()* call.

**srun does not understand mpmc instructions** Simply disable them with the flag `-DADIOS2_RUN_MPI_MPMC_TESTS=OFF`

**Tests timeout** Since we launch every tests with srun the scheduling times can exceed the test default timeout. Use a large timeout (5mins) for running your tests.

Examples of launching ADIOS2 SST unit tests using MPI DP:

```
# We omit some of the srun (SLURM) arguments which are specific of the project
# you are working on. Note that you could avoid calling srun directly by
# setting the CMAKE variable `MPIEXEC_EXECUTABLE`.

# Launch simple writer test instance
srun {PROJFLAGS} -N 1 /gpfs/alpine/proj-shared/csc331/vbolea/ADIOS2-build/bin/
↪TestCommonWrite SST mpi_dp_test CCommPattern=Min,MarshalMethod=BP5

# On another terminal launch multiple instances of the Reader test
srun {PROJFLAGS} -N 2 /gpfs/alpine/proj-shared/csc331/vbolea/ADIOS2-build/bin/
↪TestCommonRead SST mpi_dp_test
```

Alternatively, you can configure your CMake build to use srun directly:

```
cmake . -DMPIEXEC_EXECUTABLE:FILEPATH="/usr/bin/srun" \
-DMPIEXEC_EXTRA_FLAGS:STRING="-A{YourProject} -pbatch -t10" \
-DMPIEXEC_NUMPROC_FLAG:STRING="-N" \
-DMPIEXEC_MAX_NUMPROCS:STRING="-8" \
```

(continues on next page)

(continued from previous page)

```
-DADIOS2_RUN_MPI_MPMD_TESTS=OFF  
  
cmake --build .  
ctest  
  
# monitor your jobs  
watch -n1 squeue -l -u $USER
```

## OVERVIEW

In this tutorial we will learn about how to build ADIOS2, and go through several tutorials explaining basic topics.

More specifically, we will go through the following examples:

1. *Download And Build*
2. Basic tutorials:
  1. *Hello World*
  2. *Array Variables*
  3. *Attributes*
  4. *Operators*
  5. *Steps*

More advance tutorials that cover more information related to:

- Running ADIOS2 at scale (through files or streaming) with hands-on excercises: [Exascale I/O tutorial](#)
- Using ADIOS2 with Paraview, TAU, Catalyst, FIDES, VTK-M: [ADIOS2 tutorial at SC23](#)



## DOWNLOAD AND BUILD

First, you need to clone the ADIOS2 repository. You can do this by running the following command:

```
git clone https://github.com/ornladios/ADIOS2.git ADIOS2
```

---

**Note:** ADIOS2 uses [CMake](#) for building, testing and installing the library and utilities. So you need to have CMake installed on your system.

---

Then, create a build directory, run CMake, and build ADIOS2:

```
cd ADIOS2
mkdir build
cd build
cmake -DADIOS2_USE_MPI=ON ..
cmake --build .
```

---

**Note:** If you want to know more about the ADIOS2's CMake options, see section [CMake Options](#).

---

All the tutorials that we will explore are existing ADIOS2 examples located in the `ADIOS2/examples` directory.

To build any of the examples, e.g. the `helloWorld` example, you can run the following commands:

```
cd Path-To-ADIOS2/examples/hello/helloWorld
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
```





## BASIC TUTORIALS

### 24.1 Hello World

Like in any language, the first program you write is the “Hello World” program.

In this tutorial, we will see how to write “Hello World from ADIOS2” and read it back with ADIOS2. So let’s dig in! Start editing the skeleton file `ADIOS2/examples/hello/helloWorld/hello-world_tutorialSkeleton.cpp`.

1. We create an ADIOS instance, and define the greeting message in our main function as follows:

```
int main()
{
    adios2::ADIOS adios();
    const std::string greeting("Hello World from ADIOS2");
    ...
    return 0;
}
```

2. Then we create a writer function in which we pass the adios instance, and the greeting message as follows:

```
void writer(adios2::ADIOS& adios, const std::string& greeting)
{
    ...
}
```

3. In this writer function, we define an IO object, a string variable for the message as follows:

```
adios2::IO io = adios.DeclareIO("hello-world-writer");
adios2::Variable<std::string> varGreeting = io.DefineVariable<std::string>("Greeting");
```

---

**Note:** Using the IO object, we can define the engine type that we want to utilize using the `io.SetEngine()` function. If `SetEngine()` is not used, the default engine type is `BPFile` which is an alias for the latest version of the BP engine of the ADIOS2 library. See [Available Engines](#) and [Supported Engines](#) for more information. It’s important to note that the file extension of an output file, although it’s not a good practice, it can differ from the engine type, e.g. write a `foo.h5` file with the `BPFile` engine. When reading `foo.h5` you should explicitly specify the engine type as `BPFile` to read it properly.

---

4. Then we open a file with the name “hello-world-cpp.bp” and write the greeting message to it as follows:

```
adios2::Engine writer = io.Open("hello-world-cpp.bp", adios2::Mode::Write);
writer.BeginStep();
writer.Put(varGreeting, greeting);
writer.EndStep();
writer.Close();
```

---

**Note:** The BeginStep and EndStep calls are optional when **writing** one step, but they are required for multiple steps, so it is a good practice to always use them.

---

5. Now we create a reader function in which we pass the adios instance, and get the greeting message back as follows:

```
std::string reader(adios2::ADIOS& adios)
{
    ...
}
```

6. In this reader function, we define an IO object and inquire a string variable for the message as follows:

```
adios2::IO io = adios.DeclareIO("hello-world-reader");
reader.BeginStep();
adios2::Variable<std::string> varGreeting = io.InquireVariable<std::string>("Greeting");
```

7. Then we open the file with the name “hello-world-cpp.bp”, read the greeting message from it and return it as follows:

```
adios2::Engine reader = io.Open("hello-world-cpp.bp", adios2::Mode::Read);
std::string greeting;
reader.BeginStep();
reader.Get(varGreeting, greeting);
reader.EndStep();
reader.Close();
return greeting;
```

---

**Note:** In Mode::Read, the BeginStep and EndStep calls are required when **reading** one step and multiple steps. We will see in another tutorial how to read multiple steps. It’s important to note that the BeginStep should be called **before** all Inquire\* / Available\* function calls.

---

8. Finally, we call the writer and reader functions in our main function as follows:

```
int main()
{
    adios2::ADIOS adios();
    const std::string greeting("Hello World from ADIOS2");
    writer(adios, greeting);
    std::string message = reader(adios);
    std::cout << message << std::endl;
    return 0;
}
```

9. The final code should look as follows (excluding try/catch and the optional usage of MPI), and it was derived from the example `ADIOS2/examples/hello/helloWorld/hello-world.cpp`.

```

/*
 * Distributed under the OSI-approved Apache License, Version 2.0.  See
 * accompanying file Copyright.txt for details.
 *
 * hello-world.cpp : adios2 low-level API example to write and read a
 *                   std::string Variable with a greeting
 *
 * Created on: Nov 14, 2019
 * Author: William F Godoy godoywf@ornl.gov
 */

#include <iostream>
#include <stdexcept>

#include <adios2.h>
#if ADIOS2_USE_MPI
#include <mpi.h>
#endif

void writer(adios2::ADIOS &adios, const std::string &greeting)
{
    adios2::IO io = adios.DeclareIO("hello-world-writer");
    adios2::Variable<std::string> varGreeting = io.DefineVariable<std::string>("Greeting
↪");

    adios2::Engine writer = io.Open("hello-world-cpp.bp", adios2::Mode::Write);
    writer.BeginStep();
    writer.Put(varGreeting, greeting);
    writer.EndStep();
    writer.Close();
}

std::string reader(adios2::ADIOS &adios)
{
    adios2::IO io = adios.DeclareIO("hello-world-reader");
    adios2::Engine reader = io.Open("hello-world-cpp.bp", adios2::Mode::Read);
    reader.BeginStep();
    adios2::Variable<std::string> varGreeting = io.InquireVariable<std::string>("Greeting
↪");
    std::string greeting;
    reader.Get(varGreeting, greeting);
    reader.EndStep();
    reader.Close();
    return greeting;
}

int main(int argc, char *argv[])
{
    #if ADIOS2_USE_MPI
    MPI_Init(&argc, &argv);

```

(continues on next page)

(continued from previous page)

```

#endif

    try
    {
#ifdef ADIOS2_USE_MPI
        adios2::ADIOS adios(MPI_COMM_WORLD);
#else
        adios2::ADIOS adios;
#endif

        const std::string greeting = "Hello World from ADIOS2";
        writer(adios, greeting);

        const std::string message = reader(adios);
        std::cout << message << "\n";
    }
    catch (std::exception &e)
    {
        std::cout << "ERROR: ADIOS2 exception: " << e.what() << "\n";
#ifdef ADIOS2_USE_MPI
        MPI_Abort(MPI_COMM_WORLD, -1);
#endif
    }

#ifdef ADIOS2_USE_MPI
    MPI_Finalize();
#endif

    return 0;
}

```

10. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/helloWorld
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
./adios2_hello_helloWorld

```

11. You can check the content of the output file “hello-world-cpp.bp” using *bpls* as follows:

```

Path-To-ADIOS2/build/bin/bpls ./hello-world-cpp.bp

string    Greeting    scalar

```

12. The Python version of this tutorial can be found at [ADIOS2/examples/hello/helloWorld/hello-world.py](#). and it looks as follows:

```

#
# Distributed under the OSI-approved Apache License, Version 2.0.  See
# accompanying file Copyright.txt for details.
#

```

(continues on next page)

(continued from previous page)

```

# hello-world.py : adios2 Python API example to write and read a
#                 string Variable with a greeting
#
# Created on: 2/2/2021
# Author: Dmitry Ganyushin ganyushindi@ornl.gov
#
import sys
from mpi4py import MPI
from adios2 import Stream, FileReader

DATA_FILENAME = "hello-world-py.bp"
# MPI
comm = MPI.COMM_WORLD
rank = comm.Get_rank()
size = comm.Get_size()

def writer(greeting):
    """write a string to a bp file"""
    with Stream(DATA_FILENAME, "w", comm) as fh:
        fh.write("Greeting", greeting)
    return 0

def reader():
    """read a string from a bp file as Stream"""
    message = f"variable Greeting not found in {DATA_FILENAME}"
    with Stream(DATA_FILENAME, "r", comm) as fh:
        for _ in fh.steps():
            message = fh.read("Greeting")
    return message

def filereader():
    """read a string from a bp file using random access read mode"""
    with FileReader(DATA_FILENAME, comm) as fh:
        message = fh.read("Greeting")
    return message

def main():
    """driver function"""
    greeting = "Hello World from ADIOS2"
    writer(greeting)
    message = reader()
    print("As read from adios2.Stream: {}".format(message))
    message2 = filereader()
    print("As read from adios2.FileReader: {}".format(message2))
    return 0

if __name__ == "__main__":

```

(continues on next page)

```
sys.exit(main())
```

## 24.2 Variables

In the previous tutorial we learned how to define a simple string variable, write it, and read it back.

In this tutorial we will go two steps further:

1. We will define variables which include arrays, and we will write them and read them back.
2. We will use MPI to write and read the above variables in parallel.

Let's start with the writing part.

Start editing the skeleton file `ADIOS2/examples/hello/bpWriter/bpWriter_tutorialSkeleton.cpp`.

1. In an MPI application first we need to always initialize MPI. We do that with the following lines:

```
int rank, size;
int provided;

// MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);
```

2. Now we need to create some application variables which will be used to define ADIOS2 variables.

```
// Application variable
std::vector<float> myFloats = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9};
std::vector<int> myInts = {0, -1, -2, -3, -4, -5, -6, -7, -8, -9};
const std::size_t Nx = myFloats.size();
const std::string myString("Hello Variable String from rank " + std::to_string(rank));
```

3. Now we need to define an ADIOS2 instance and the ADIOS2 variables.

```
adios2::ADIOS adios(MPI_COMM_WORLD);
adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

adios2::Variable<float> bpFloats = bpIO.DefineVariable<float>(
    "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);

adios2::Variable<int> bpInts = bpIO.DefineVariable<int>("bpInts", {size * Nx}, {rank *
    Nx}, {Nx}, adios2::ConstantDims);

// For the sake of the tutorial we create an unused variable
adios2::Variable<std::string> bpString = bpIO.DefineVariable<std::string>("bpString");
```

**Note:** The above int/float variables are global arrays. The 1st argument of the `DefineVariable` function is the variable name, the 2nd are the global dimensions, the 3rd is the start index for a rank, the 4th are the rank/local

dimensions, and the 5th is a boolean variable to indicate if the dimensions are constant or not over multiple steps, where `adios2::ConstantDims == true`. We will explore other tutorials that don't use constant dimensions.

4. Now we need to open the ADIOS2 engine and write the variables.

```
adios2::Engine bpWriter = bpIO.Open("myVector_cpp.bp", adios2::Mode::Write);

bpWriter.BeginStep();
bpWriter.Put(bpFloats, myFloats.data());
bpWriter.Put(bpInts, myInts.data());
// bpWriter.Put(bpString, myString);
bpWriter.EndStep();

bpWriter.Close();
```

5. Finally we need to finalize MPI.

```
MPI_Finalize();
```

6. The final code should look as follows (excluding try/catch and the optional usage of MPI), and it was derived from the example `ADIOS2/examples/hello/bpWriter/bpWriter.cpp`.

```
/*
 * Distributed under the OSI-approved Apache License, Version 2.0. See
 * accompanying file Copyright.txt for details.
 *
 * bpWriter.cpp: Simple self-descriptive example of how to write a variable
 * to a BP File that lives in several MPI processes.
 *
 * Created on: Feb 16, 2017
 * Author: William F Godoy godoywf@ornl.gov
 */

#include <ios>           //std::ios_base::failure
#include <iostream>      //std::cout
#include <stdexcept>     //std::invalid_argument std::exception
#include <vector>

#include <adios2.h>
#ifdef ADIOS2_USE_MPI
#include <mpi.h>
#endif

int main(int argc, char *argv[])
{
    int rank, size;
#ifdef ADIOS2_USE_MPI
    int provided;

    // MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
```

(continues on next page)

(continued from previous page)

```

#else
    rank = 0;
    size = 1;
#endif

    /** Application variable */
    std::vector<float> myFloats = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9};
    std::vector<int> myInts = {0, -1, -2, -3, -4, -5, -6, -7, -8, -9};
    const std::size_t Nx = myFloats.size();

    const std::string myString("Hello Variable String from rank " + std::to_
↪string(rank));

    try
    {
        /** ADIOS class factory of IO class objects */
    #if ADIOS2_USE_MPI
        adios2::ADIOS adios(MPI_COMM_WORLD);
    #else
        adios2::ADIOS adios;
    #endif

        /** IO class object: settings and factory of Settings: Variables,
         * Parameters, Transports, and Execution: Engines */
        adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

        /** global array : name, { shape (total) }, { start (local) }, {
         * count
         * (local) }, all are constant dimensions */
        adios2::Variable<float> bpFloats = bpIO.DefineVariable<float>(
            "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);

        adios2::Variable<int> bpInts = bpIO.DefineVariable<int>("bpInts", {size * Nx},
↪{rank * Nx},
                                                    {Nx},
↪adios2::ConstantDims);

        adios2::Variable<std::string> bpString = bpIO.DefineVariable<std::string>(
↪"bpString");
        (void)bpString; // For the sake of the example we create an unused
                        // variable

        std::string filename = "myVector_cpp.bp";
        /** Engine derived class, spawned to start IO operations */
        adios2::Engine bpWriter = bpIO.Open(filename, adios2::Mode::Write);

        bpWriter.BeginStep();
        /** Put variables for buffering, template type is optional */
        bpWriter.Put(bpFloats, myFloats.data());
        bpWriter.Put(bpInts, myInts.data());
        // bpWriter.Put(bpString, myString);
        bpWriter.EndStep();

```

(continues on next page)



(continued from previous page)

```

    /** Create bp file, engine becomes unreachable after this*/
    bpWriter.Close();
    if (rank == 0)
    {
        std::cout << "Wrote file " << filename
                    << " to disk. It can now be read by running "
                    << "./bin/adios2_hello_bpReader.\n";
    }
}
catch (std::invalid_argument &e)
{
    std::cerr << "Invalid argument exception: " << e.what() << "\n";
#if ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}
catch (std::ios_base::failure &e)
{
    std::cerr << "IO System base failure exception: " << e.what() << "\n";
#if ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}
catch (std::exception &e)
{
    std::cerr << "Exception: " << e.what() << "\n";
#if ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}

#if ADIOS2_USE_MPI
    MPI_Finalize();
#endif

    return 0;
}

```

7. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/bpWriter
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
mpirun -np 2 ./adios2_hello_bpWriter_mpi

```

8. You can check the content of the output file “myVector\_cpp.bp” using *bpls* as follows:

```
Path-To-ADIOS2/build/bin/bpls ./myVector_cpp.bp
```

```
float    bpFloats {10}
int32_t  bpInts   {10}
```

Now let's move to the reading part.

Start editing the skeleton file `ADIOS2/examples/hello/bpReader/bpReader_tutorialSkeleton.cpp`.

9. In an MPI application first we need to always initialize MPI. We do that with the following line:

```
int rank, size;
int provided;

// MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);
```

10. Now we need to define an ADIOS2 instance and open the ADIOS2 engine.

```
adios2::ADIOS adios(MPI_COMM_WORLD);

adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

adios2::Engine bpReader = bpIO.Open("myVector_cpp.bp", adios2::Mode::Read);
```

11. Now we need to read the variables. In this case we know the variables that we need to inquire, so we can use the `InquireVariable` function immediately. But let's explore how to check the available variables in a file first, and then we will use the `InquireVariable` function.

```
bpReader.BeginStep();
const std::map<std::string, adios2::Params> variables = bpIO.AvailableVariables();

for (const auto &variablePair : variables)
{
    std::cout << "Name: " << variablePair.first;
    for (const auto &parameter : variablePair.second)
    {
        std::cout << "\t" << parameter.first << ": " << parameter.second << "\n";
    }
}

adios2::Variable<float> bpFloats = bpIO.InquireVariable<float>("bpFloats");
adios2::Variable<int> bpInts = bpIO.InquireVariable<int>("bpInts");
```

12. Now we need to read the variables from each rank. We will use the `SetSelection` to set the start index and rank dimensions, then `Get` function to read the variables, and print the contents from rank 0.

```
const std::size_t Nx = 10;
if (bpFloats) // means found
{
    std::vector<float> myFloats;
```

(continues on next page)

(continued from previous page)

```

// read only the chunk corresponding to our rank
bpFloats.SetSelection({{Nx * rank}, {Nx}});
bpReader.Get(bpFloats, myFloats, adios2::Mode::Sync);

if (rank == 0)
{
    std::cout << "MyFloats: \n";
    for (const auto number : myFloats)
    {
        std::cout << number << " ";
    }
    std::cout << "\n";
}

if (bpInts) // means not found
{
    std::vector<int> myInts;
    // read only the chunk corresponding to our rank
    bpInts.SetSelection({{Nx * rank}, {Nx}});

    bpReader.Get(bpInts, myInts, adios2::Mode::Sync);

    if (rank == 0)
    {
        std::cout << "myInts: \n";
        for (const auto number : myInts)
        {
            std::cout << number << " ";
        }
        std::cout << "\n";
    }
}

```

**Note:** While using the Get function, we used the third parameter named Mode. The mode parameter can also be used for the Put function.

For the Put function, there are three modes: Deferred (default), Sync, and Span. and for the Get there are two modes: Deferred (default) and Sync.

1. The Deferred mode is the default mode, because it is the fastest mode, as it allows Put / Get to be grouped before potential data transport at the first encounter of PerformPuts / PerformGets, EndStep or Close.
2. The Sync mode forces Put / Get to be performed immediately so that the data are available immediately.
3. The Span mode is special mode of Deferred that allows population from non-contiguous memory structures.

For more information about the Mode parameter for both Put and Get functions, and when you should use each option see [Basics: Interface Components: Engine](#).

13. Now we need close the ADIOS2 engine.

```
bpReader.EndStep();  
bpReader.Close();
```

14. Finally we need to finalize MPI.

```
MPI_Finalize();
```

15. The final code should look as follows (excluding try/catch), and it was derived from the example `ADIOS2/examples/hello/bpWriter/bpWriter.cpp`.

```
/*  
 * Distributed under the OSI-approved Apache License, Version 2.0.  See  
 * accompanying file Copyright.txt for details.  
 *  
 * bpReader.cpp: Simple self-descriptive example of how to read a variable  
 * from a BP File.  
 *  
 * Created on: Feb 16, 2017  
 * Author: William F Godoy godoywf@ornl.gov  
 */  
#include <ios>           //std::ios_base::failure  
#include <iostream>      //std::cout  
#if ADIOS2_USE_MPI  
#include <mpi.h>  
#endif  
#include <stdexcept>     //std::invalid_argument std::exception  
#include <vector>  
  
#include <adios2.h>  
  
int main(int argc, char *argv[])  
{  
    int rank, size;  
  
#if ADIOS2_USE_MPI  
    int provided;  
    // MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP  
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);  
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
    MPI_Comm_size(MPI_COMM_WORLD, &size);  
#else  
    rank = 0;  
    size = 1;  
#endif  
    std::cout << "rank " << rank << " size " << size << "\n";  
    try  
    {  
#if ADIOS2_USE_MPI  
        /** ADIOS class factory of IO class objects */  
        adios2::ADIOS adios(MPI_COMM_WORLD);  
#else  
        adios2::ADIOS adios;  
#endif  
    }
```

(continues on next page)

(continued from previous page)

```

/** IO class object: settings and factory of Settings: Variables,
 * Parameters, Transports, and Execution: Engines */
adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

/** Engine derived class, spawned to start IO operations */
adios2::Engine bpReader = bpIO.Open("myVector_cpp.bp", adios2::Mode::Read);

bpReader.BeginStep();
const std::map<std::string, adios2::Params> variables = bpIO.
↪AvailableVariables();

for (const auto &variablePair : variables)
{
    std::cout << "Name: " << variablePair.first;

    for (const auto &parameter : variablePair.second)
    {
        std::cout << "\t" << parameter.first << ": " << parameter.second << "\n";
    }
}

/** Write variable for buffering */
adios2::Variable<float> bpFloats = bpIO.InquireVariable<float>("bpFloats");
adios2::Variable<int> bpInts = bpIO.InquireVariable<int>("bpInts");

const std::size_t Nx = 10;
if (bpFloats) // means found
{
    std::vector<float> myFloats;

    // read only the chunk corresponding to our rank
    bpFloats.SetSelection({{Nx * rank}, {Nx}});
    // myFloats.data is pre-allocated
    bpReader.Get(bpFloats, myFloats, adios2::Mode::Sync);

    if (rank == 0)
    {
        std::cout << "MyFloats: \n";
        for (const auto number : myFloats)
        {
            std::cout << number << " ";
        }
        std::cout << "\n";
    }
}

if (bpInts) // means not found
{
    std::vector<int> myInts;
    // read only the chunk corresponding to our rank
    bpInts.SetSelection({{Nx * rank}, {Nx}});

```

(continues on next page)

(continued from previous page)

```

        bpReader.Get(bpInts, myInts, adios2::Mode::Sync);

        if (rank == 0)
        {
            std::cout << "myInts: \n";
            for (const auto number : myInts)
            {
                std::cout << number << " ";
            }
            std::cout << "\n";
        }
    }
    bpReader.EndStep();

    /** Close bp file, engine becomes unreachable after this*/
    bpReader.Close();
}
catch (std::invalid_argument &e)
{
    if (rank == 0)
    {
        std::cerr << "Invalid argument exception, STOPPING PROGRAM from rank " <<
↪rank << "\n";
        std::cerr << e.what() << "\n";
    }
    #if ADIOS2_USE_MPI
        MPI_Abort(MPI_COMM_WORLD, 1);
    #endif
}
catch (std::ios_base::failure &e)
{
    if (rank == 0)
    {
        std::cerr << "IO System base failure exception, STOPPING PROGRAM "
                    << "from rank "
                    << rank << "\n";
        std::cerr << e.what() << "\n";
        std::cerr << "The file myVector_cpp.bp does not exist."
                    << " Presumably this is because adios2_hello_bpWriter has not "
                    << "been run."
                    << " Run ./adios2_hello_bpWriter before running this program.\n";
    }
    #if ADIOS2_USE_MPI
        MPI_Abort(MPI_COMM_WORLD, 1);
    #endif
}
catch (std::exception &e)
{
    if (rank == 0)
    {
        std::cerr << "Exception, STOPPING PROGRAM from rank " << rank << "\n";
    }
}

```

(continues on next page)

(continued from previous page)

```

        std::cerr << e.what() << "\n";
    }
#ifdef ADIOS2_USE_MPI
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}

#ifdef ADIOS2_USE_MPI
    MPI_Finalize();
#endif

    return 0;
}

```

16. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/bpReader
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
mpirun -np 2 ./adios2_hello_bpReader_mpi

```

## 24.3 Attributes

In the previous tutorial, we learned how to write/read variables.

In this tutorial, we will explore how to write/read attributes. Attributes are metadata related to the whole dataset or to a specific variable. In this tutorial, we will only focus on attributes related to the whole dataset, but we will explain how variable's attributes can be used too.

Start editing the skeleton file `ADIOS2/examples/hello/bpAttributeWriteRead/bpAttributeWriteRead_tutorialSkeleton.cpp`.

1. In an MPI application first we need to always initialize MPI. We do that with the following lines:

```

int rank, size;
int provided;

// MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);

```

2. Now we need to create a application variable which will be used to define an ADIOS2 variable.

```
std::vector<float> myFloats = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9};
```

3. Then, we need to create an ADIOS2 instance.

```
adios2::ADIOS adios(MPI_COMM_WORLD);
```

4. Then, we create the following writer function:

```
void writer(adios2::ADIOS &adios, int rank, int size, std::vector<float> &myFloats)
{
    ...
}
```

5. In this writer function, we define an IO object, and a float vector variable as follows:

```
adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

const std::size_t Nx = myFloats.size();
adios2::Variable<float> bpFloats = bpIO.DefineVariable<float>(
    "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);
```

6. Now, we will define various types of attributes as follows:

```
bpIO.DefineAttribute<std::string>("Single_String", "File generated with ADIOS2");

std::vector<std::string> myStrings = {"one", "two", "three"};
bpIO.DefineAttribute<std::string>("Array_of_Strings", myStrings.data(), myStrings.
    ↪size());

bpIO.DefineAttribute<double>("Attr_Double", 0.f);
std::vector<double> myDoubles = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10};
bpIO.DefineAttribute<double>("Array_of_Doubles", myDoubles.data(), myDoubles.size());
```

---

**Note:** if we want to define an attribute for a specific variable, we can use one of the following API:

```
template <class T>
Attribute<T> DefineAttribute(const std::string &name, const T *data, const size_t size,
    const std::string &variableName = "", const std::string_
    ↪separator = "/",
    const bool allowModification = false);

template <class T>
Attribute<T> DefineAttribute(const std::string &name, const T &value,
    const std::string &variableName = "", const std::string_
    ↪separator = "/",
    const bool allowModification = false);
```

As we can see, by default the attributes don't change over multiple steps, but we can change that by setting `allowModification` to true.

---

7. Then, we open a file for writing:

```
adios2::Engine bpWriter = bpIO.Open("fileAttributes.bp", adios2::Mode::Write);
```

8. Now, we write the data and close the file:

```
bpWriter.BeginStep();
bpWriter.Put<float>(bpFloats, myFloats.data());
bpWriter.EndStep();
bpWriter.Close();
```



9. Steps 1-8 are used for writing, we will define a reader function in the rest of the steps:

```
void reader(adios2::ADIOS &adios, int rank, int size)
{
    ...
}
```

10. In this reader function, we define an IO object, and open the file for reading:

```
adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");
adios2::Engine bpReader = bpIO.Open("fileAttributes.bp", adios2::Mode::Read);
```

11. Now, we check the available attributes as follows:

```
bpReader.BeginStep();
const auto attributesInfo = bpIO.AvailableAttributes();

for (const auto &attributeInfoPair : attributesInfo)
{
    std::cout << "Attribute: " << attributeInfoPair.first;
    for (const auto &attributePair : attributeInfoPair.second)
    {
        std::cout << "\tKey: " << attributePair.first << "\tValue: " << attributePair.
↪second
                << "\n";
    }
    std::cout << "\n";
}
```

12. Now we will inquire and get the attributes as follows:

```
adios2::Attribute<float> singleString = bpIO.InquireAttribute<float>("Single_String");
if (singleString)
{
    std::cout << singleString.Name() << ": " << singleString.Data()[0] << "\n";
}
adios2::Attribute<std::string> arrayOfStrings =
    bpIO.InquireAttribute<std::string>("Array_of_Strings");
if (arrayOfStrings)
{
    std::cout << arrayOfStrings.Name() << ": ";
    for (const auto &value : arrayOfStrings.Data())
    {
        std::cout << value << " ";
    }
    std::cout << "\n";
}
adios2::Attribute<double> attrDouble = bpIO.InquireAttribute<double>("Attr_Double");
if (attrDouble)
{
    std::cout << attrDouble.Name() << ": " << attrDouble.Data()[0] << "\n";
}
adios2::Attribute<double> arrayOfDoubles = bpIO.InquireAttribute<double>("Array_of_
↪Doubles");
```

(continues on next page)

(continued from previous page)

```

if (arrayOfDoubles)
{
    std::cout << arrayOfDoubles.Name() << ": ";
    for (const auto &value : arrayOfDoubles.Data())
    {
        std::cout << value << " ";
    }
    std::cout << "\n";
}

```

13. Afterward, we will inquire and get the variable as follows:

```

adios2::Variable<float> bpFloats = bpIO.InquireVariable<float>("bpFloats");
const std::size_t Nx = 10;
std::vector<float> myFloats(Nx);
if (bpFloats)
{
    bpFloats.SetSelection({{Nx * rank}, {Nx}});
    bpReader.Get(bpFloats, myFloats.data());
}
bpReader.EndStep();

```

14. Finally, we close the file:

```
bpReader.Close();
```

15. In the main function, we call the writer and reader functions as follows:

```

writer(adios, rank, size, myFloats);
reader(adios, rank, size);

```

16. Finally, we finalize MPI:

```
MPI_Finalize();
```

17. The final code should look as follows (excluding try/catch and optional usage MPI), and it was derived from the example `ADIOS2/examples/hello/bpAttributeWriteRead/bpAttributeWriteRead.cpp`.

```

/*
 * Distributed under the OSI-approved Apache License, Version 2.0.  See
 * accompanying file Copyright.txt for details.
 *
 * bpAttributeWriteRead.cpp: Simple self-descriptive example of how to write/read
 * attributes and
 * a variable to a BP File that lives in several MPI processes.
 *
 * Created on: Feb 16, 2017
 * Author: William F Godoy godoywf@ornl.gov
 */

#include <ios>           //std::ios_base::failure
#include <iostream>      //std::cout
#if ADIOS2_USE_MPI

```

(continues on next page)

(continued from previous page)

```

#include <mpi.h>
#endif
#include <stdexcept> //std::invalid_argument std::exception
#include <string>
#include <vector>

#include <adios2.h>

void writer(adios2::ADIOS &adios, int rank, int size, std::vector<float> &myFloats)
{
    /** IO class object: settings and factory of Settings: Variables,
     * Parameters, Transports, and Execution: Engines */
    adios2::IO bpIO = adios.DeclareIO("BPFile_N2N");

    const std::size_t Nx = myFloats.size();

    /** global array : name, { shape (total) }, { start (local) }, { count
     * (local) }, all are constant dimensions */
    adios2::Variable<float> bpFloats = bpIO.DefineVariable<float>(
        "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);

    bpIO.DefineAttribute<std::string>("Single_String", "File generated with ADIOS2");

    std::vector<std::string> myStrings = {"one", "two", "three"};
    bpIO.DefineAttribute<std::string>("Array_of_Strings", myStrings.data(), myStrings.
↪size());

    bpIO.DefineAttribute<double>("Attr_Double", 0.f);
    std::vector<double> myDoubles = {1, 2, 3, 4, 5, 6, 7, 8, 9, 10};
    bpIO.DefineAttribute<double>("Array_of_Doubles", myDoubles.data(), myDoubles.size());

    /** Engine derived class, spawned to start IO operations */
    adios2::Engine bpWriter = bpIO.Open("fileAttributes.bp", adios2::Mode::Write);

    bpWriter.BeginStep();

    /** Write variable for buffering */
    bpWriter.Put<float>(bpFloats, myFloats.data());

    bpWriter.EndStep();

    /** Create bp file, engine becomes unreachable after this*/
    bpWriter.Close();
}

void reader(adios2::ADIOS &adios, int rank, int /*size*/)
{
    adios2::IO bpIO = adios.DeclareIO("BPReader");

    adios2::Engine bpReader = bpIO.Open("fileAttributes.bp", adios2::Mode::Read);

    bpReader.BeginStep();

```

(continues on next page)

(continued from previous page)

```

const auto attributesInfo = bpIO.AvailableAttributes();

for (const auto &attributeInfoPair : attributesInfo)
{
    std::cout << "Attribute: " << attributeInfoPair.first;
    for (const auto &attributePair : attributeInfoPair.second)
    {
        std::cout << "\tKey: " << attributePair.first << "\tValue: " <<
↪ attributePair.second
        << "\n";
    }
    std::cout << "\n";
}

adios2::Attribute<float> singleString = bpIO.InquireAttribute<float>("Single_String
↪");
if (singleString)
{
    std::cout << singleString.Name() << ": " << singleString.Data()[0] << "\n";
}
adios2::Attribute<std::string> arrayOfStrings =
    bpIO.InquireAttribute<std::string>("Array_of_Strings");
if (arrayOfStrings)
{
    std::cout << arrayOfStrings.Name() << ": ";
    for (const auto &value : arrayOfStrings.Data())
    {
        std::cout << value << " ";
    }
    std::cout << "\n";
}
adios2::Attribute<double> attrDouble = bpIO.InquireAttribute<double>("Attr_Double");
if (attrDouble)
{
    std::cout << attrDouble.Name() << ": " << attrDouble.Data()[0] << "\n";
}
adios2::Attribute<double> arrayOfDoubles = bpIO.InquireAttribute<double>("Array_of_
↪Doubles");
if (arrayOfDoubles)
{
    std::cout << arrayOfDoubles.Name() << ": ";
    for (const auto &value : arrayOfDoubles.Data())
    {
        std::cout << value << " ";
    }
    std::cout << "\n";
}

adios2::Variable<float> bpFloats = bpIO.InquireVariable<float>("bpFloats");
const std::size_t Nx = 10;
std::vector<float> myFloats(Nx);
if (bpFloats)

```

(continues on next page)

(continued from previous page)

```

    {
        bpFloats.SetSelection({{Nx * rank}, {Nx}});
        bpReader.Get(bpFloats, myFloats.data());
    }

    bpReader.EndStep();

    bpReader.Close();
}

int main(int argc, char *argv[])
{
    int rank, size;
#ifdef ADIOS2_USE_MPI
    int provided;

    // MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
#else
    rank = 0;
    size = 1;
#endif

    /** Application variable */
    std::vector<float> myFloats = {0, 1, 2, 3, 4, 5, 6, 7, 8, 9};
    try
    {
        /** ADIOS class factory of IO class objects */
#ifdef ADIOS2_USE_MPI
        adios2::ADIOS adios(MPI_COMM_WORLD);
#else
        adios2::ADIOS adios;
#endif

        writer(adios, rank, size, myFloats);
        reader(adios, rank, size);
    }
    catch (std::invalid_argument &e)
    {
        std::cout << "Invalid argument exception, STOPPING PROGRAM from rank " << rank <<
        "\n";
        std::cout << e.what() << "\n";
#ifdef ADIOS2_USE_MPI
        std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
        MPI_Abort(MPI_COMM_WORLD, 1);
#endif
    }
    catch (std::ios_base::failure &e)
    {
        std::cout << "IO System base failure exception, STOPPING PROGRAM from rank " <<

```

(continues on next page)

(continued from previous page)

```

↪rank
    << "\n";
    std::cout << e.what() << "\n";
#if ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}
catch (std::exception &e)
{
    std::cout << "Exception, STOPPING PROGRAM from rank " << rank << "\n";
    std::cout << e.what() << "\n";
#if ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}

#if ADIOS2_USE_MPI
    MPI_Finalize();
#endif

    return 0;
}

```

18. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/bpAttributeWriteRead
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
mpirun -np 2 ./adios2_hello_bpAttributeWriteRead_mpi

```

19. You can check the content of the output file “fileAttributes.bp” using *bpls* as follows:

```

Path-To-ADIOS2/build/bin/bpls ./fileAttributes.bp

float    bpFloats    {20}

```

## 24.4 Operators

In the previous tutorial we learned how to write and read attributes.

For this example to work, you would need to have the SZ compression library installed, which ADIOS automatically detects. The easiest way to install SZ is with Spack, and you can do that as follows:

```

git clone https://github.com/spack/spack.git ~/spack
cd ~/spack
. share/spack/setup-env.sh

```

(continues on next page)

(continued from previous page)

```
spack install sz
spack load sz
```

In this tutorial we will learn how to use operators. Operators are used for Data compression/decompression, lossy and lossless. They act upon the user application data, either from a variable or a set of variables in a IO object.

Additionally, we will explore how to simply write variables across multiple steps.

So, let's dig in!

Start editing the skeleton file `ADIOS2/examples/hello/bpOperatorSZWriter/bpOperatorSZWriter_tutorialSkeleton.cpp`.

1. In an MPI application first we need to always initialize MPI. We do that with the following lines:

```
int rank, size;
int rank, size;
int provided;

// MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);
```

2. This application has command line arguments for the size of the data, and the compression accuracy, which we can read as follows:

```
const std::size_t Nx = static_cast<std::size_t>(std::stoull(argv[1]));
const double accuracy = std::stod(argv[2]);
```

3. Now we need to create some application variables which will be used to define ADIOS2 variables.

```
std::vector<float> myFloats(Nx);
std::vector<double> myDoubles(Nx);
std::iota(myFloats.begin(), myFloats.end(), 0.);
std::iota(myDoubles.begin(), myDoubles.end(), 0.);
```

4. Now we need to create an ADIOS2 instance and IO object.

```
adios2::ADIOS adios(MPI_COMM_WORLD);
adios2::IO bpIO = adios.DeclareIO("BPFile_SZ");
```

5. Now we need to define the variables we want to write.

```
adios2::Variable<float> bpFloats = bpIO.DefineVariable<float>(
    "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);
adios2::Variable<double> bpDoubles = bpIO.DefineVariable<double>(
    "bpDoubles", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);
```

6. Now we need to define the compression operator we want to use. In this case we will use the SZ compressor.

```
adios2::Operator op = bpIO.DefineOperator("SZCompressor", "sz");
varFloats.AddOperation(op, {"accuracy", std::to_string(accuracy)});
varDoubles.AddOperation(op, {"accuracy", std::to_string(accuracy)});
```

**Note:** `DefineOperator()`'s second parameter can be either `zfp` or `sz`. For more information regarding operators and their properties you can look at *Basics: Interface Components: Operator*.

7. Let's also create an attribute to store the accuracy value.

```
adios2::Attribute<double> attribute = bpIO.DefineAttribute<double>("accuracy", accuracy);
```

8. Now we need to open the file for writing.

```
adios2::Engine bpWriter = bpIO.Open("SZexample.bp", adios2::Mode::Write);
```

9. Now we need to write the data. We will write the data for 3 steps, and edit them in between.

```
for (unsigned int step = 0; step < 3; ++step)
{
    bpWriter.BeginStep();

    bpWriter.Put<double>(bpDoubles, myDoubles.data());
    bpWriter.Put<float>(bpFloats, myFloats.data());

    bpWriter.EndStep();

    // here you can modify myFloats, myDoubles per step
    std::transform(myFloats.begin(), myFloats.end(), myFloats.begin(),
        [&](float v) -> float { return 2 * v; });
    std::transform(myDoubles.begin(), myDoubles.end(), myDoubles.begin(),
        [&](double v) -> double { return 3 * v; });
}
```

10. Now we need to close the file.

```
bpWriter.Close();
```

11. Finally we need to finalize MPI.

```
MPI_Finalize();
```

12. The final code should look as follows (excluding try/catch and optional usage of MPI), and it was derived from the example `ADIOS2/examples/hello/bpOperatorSZWriter/bpOperatorSZWriter.cpp`.

```
/*
 * Distributed under the OSI-approved Apache License, Version 2.0. See
 * accompanying file Copyright.txt for details.
 *
 * bpOperatorSZWriter.cpp : example using operator by passing compression arguments
 *
 * Created on: Aug 3, 2018
 * Author: William F Godoy godoywf@ornl.gov
 */

#include <algorithm> //std::transform
#include <ios>       //std::ios_base::failure
#include <iostream>  //std::cout
```

(continues on next page)



(continued from previous page)

```

#include <numeric>    //std::iota
#include <stdexcept>  //std::invalid_argument std::exception
#include <vector>

#include "adios2.h"
#if ADIOS2_USE_MPI
#include <mpi.h>
#endif

void Usage()
{
    std::cout << "\n";
    std::cout << "USAGE:\n";
    std::cout << "./adios2_hello_bpOperatorSZWriter Nx sz_accuracy\n";
    std::cout << "\t Nx: size of float and double arrays to be compressed\n";
    std::cout << "\t sz_accuracy: absolute accuracy e.g. 0.1, 0.001, to skip "
                "compression: -1\n\n";
}

int main(int argc, char *argv[])
{
    if (argc != 3)
    {
        Usage();
        return EXIT_SUCCESS;
    }

    int rank, size;
#if ADIOS2_USE_MPI
    int provided;

    // MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
#else
    rank = 0;
    size = 1;
#endif

    try
    {
        const std::size_t Nx = static_cast<std::size_t>(std::stoull(argv[1]));
        const double accuracy = std::stod(argv[2]);

        /** Application variable */
        std::vector<float> myFloats(Nx);
        std::vector<double> myDoubles(Nx);
        std::iota(myFloats.begin(), myFloats.end(), 0.);
        std::iota(myDoubles.begin(), myDoubles.end(), 0.);

        /** ADIOS class factory of IO class objects */

```

(continues on next page)

(continued from previous page)

```

#if ADIOS2_USE_MPI
    adios2::ADIOS adios(MPI_COMM_WORLD);
#else
    adios2::ADIOS adios;
#endif

    /** IO class object: settings and factory of Settings: Variables,
     * Parameters, Transports, and Execution: Engines */
    adios2::IO bpIO = adios.DeclareIO("BPFile_SZ");

    adios2::Variable<float> varFloats = bpIO.DefineVariable<float>(
        "bpFloats", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);

    adios2::Variable<double> varDoubles = bpIO.DefineVariable<double>(
        "bpDoubles", {size * Nx}, {rank * Nx}, {Nx}, adios2::ConstantDims);

    if (accuracy > 1E-16)
    {
        adios2::Operator op = adios.DefineOperator("SZCompressor", "sz");
        varFloats.AddOperation(op, {"accuracy", std::to_string(accuracy)});
        varDoubles.AddOperation(op, {"accuracy", std::to_string(accuracy)});
    }

    adios2::Attribute<double> attribute = bpIO.DefineAttribute<double>("SZ_accuracy",
↪ accuracy);

    // To avoid compiling warnings
    (void)attribute;

    /** Engine derived class, spawned to start IO operations */
    adios2::Engine bpWriter = bpIO.Open("SZexample.bp", adios2::Mode::Write);

    for (unsigned int step = 0; step < 3; ++step)
    {
        bpWriter.BeginStep();

        bpWriter.Put(varFloats, myFloats.data());
        bpWriter.Put(varDoubles, myDoubles.data());

        bpWriter.EndStep();

        // here you can modify myFloats, myDoubles per step
        std::transform(myFloats.begin(), myFloats.end(), myFloats.begin(),
            [&](float v) -> float { return 2 * v; });
        std::transform(myDoubles.begin(), myDoubles.end(), myDoubles.begin(),
            [&](double v) -> double { return 3 * v; });
    }

    /** Create bp file, engine becomes unreachable after this*/
    bpWriter.Close();
}
catch (std::invalid_argument &e)

```

(continues on next page)

(continued from previous page)

```

{
    std::cerr << "Invalid argument exception: " << e.what() << "\n";
#ifdef ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}
catch (std::ios_base::failure &e)
{
    std::cerr << "IO System base failure exception: " << e.what() << "\n";
#ifdef ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}
catch (std::exception &e)
{
    std::cerr << "Exception: " << e.what() << "\n";
#ifdef ADIOS2_USE_MPI
    std::cerr << "STOPPING PROGRAM from rank " << rank << "\n";
    MPI_Abort(MPI_COMM_WORLD, 1);
#endif
}

#ifdef ADIOS2_USE_MPI
    MPI_Finalize();
#endif

    return 0;
}

```

13. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/bpOperatorSZWriter
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
mpirun -np 2 ./adios2_hello_bpOperatorSZWriter_mpi 20 0.000001

```

12. You can check the content of the output file “SZexample.bp” using *bpls* as follows:

```

Path-To-ADIOS2/build/bin/bpls ./SZexample.bp

double    bpDoubles    3*{40}
float     bpFloats     3*{40}

```

## 24.5 Steps

In the previous tutorial, we introduced the concept of operators, and briefly touched upon the concept of steps.

In this tutorial, we will explore how to write data for multiple steps, and how to read them back.

So let's dig in!

Start editing the skeleton file `ADIOS2/examples/hello/bpStepsWriteRead/bpStepsWriteRead_tutorialSkeleton.cpp`.

1. In an MPI application first we need to always initialize MPI. We do that with the following lines:

```
int rank, size;
int rank, size;
int provided;

// MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
MPI_Comm_size(MPI_COMM_WORLD, &size);
```

2. This application has an optional command line argument for engine being used. If no argument is provided, the default engine is BPFile.

```
const std::string engine = argv[1] ? argv[1] : "BPFile";
```

3. We will define the number of steps and the size of the data that we will create.

```
const std::string filename = engine + "StepsWriteRead.bp";
const unsigned int nSteps = 10;
const unsigned int Nx = 60000;
```

4. Now we need to create an ADIOS2 instance.

```
adios2::ADIOS adios(MPI_COMM_WORLD);
```

5. Now we will populate the writer function with the following signature:

```
void writer(adios2::ADIOS &adios, const std::string &engine, const std::string &fname,
            const size_t Nx, unsigned int nSteps, int rank, int size)
{
    ...
}
```

6. Let's create some simulation data. We will create a 1D array of size Nx, and fill it with 0.block

```
std::vector<double> simData(Nx, 0.0);
```

7. Now we will create an IO object and set the engine type.block

```
adios2::IO bpIO = adios.DeclareIO("SimulationOutput");
io.SetEngine(engine);
```

---

**Note:** The beauty of ADIOS2 is that you write the same code for all engines. The only thing that changes is the engine name. The underlying engine handles all the intricacies of the engine's format, and the user enjoys the API's simplicity.

---

8. Now we will create a variable for the simulation data and the step.

```
const adios2::Dims shape{static_cast<size_t>(size * Nx)};
const adios2::Dims start{static_cast<size_t>(rank * Nx)};
const adios2::Dims count{Nx};
auto bpFloats = bpIO.DefineVariable<float>("bpFloats", shape, start, count);

auto bpStep = bpIO.DefineVariable<unsigned int>("bpStep");
```

9. Now we will open the file for writing.

```
adios2::Engine bpWriter = bpIO.Open(fname, adios2::Mode::Write);
```

10. Now we will write the data for each step.

```
for (unsigned int step = 0; step < nSteps; ++step)
{
    const adios2::Box<adios2::Dims> sel({0}, {Nx});
    bpFloats.SetSelection(sel);

    bpWriter.BeginStep();
    bpWriter.Put(bpFloats, simData.data());
    bpWriter.Put(bpStep, step);
    bpWriter.EndStep();

    // Update values in the simulation data
    update_array(simData, 10);
}
```

11. Now we will close the file.

```
bpWriter.Close();
```

12. Now we will populate the reader function with the following signature:

```
void reader(adios2::ADIOS &adios, const std::string &engine, const std::string &fname,
            const size_t Nx, unsigned int /*nSteps*/, int rank, int /*size*/)
{
    ...
}
```

13. Now we will create an IO object and set the engine type.

```
adios2::IO bpIO = adios.DeclareIO("SimulationOutput");
io.SetEngine(engine);
```

14. Now we will open the file for reading.

```
adios2::Engine bpReader = bpIO.Open(fname, adios2::Mode::Read);
```

15. Now we will create a vector to store simData and a variable for the step.

```
std::vector<float> simData(Nx, 0);
unsigned int inStep = 0;
```

16. Now we will read the data for each step.

```

for (unsigned int step = 0; bpReader.BeginStep() == adios2::StepStatus::OK; ++step)
{
    auto bpFloats = bpIO.InquireVariable<float>("bpFloats");
    if (bpFloats)
    {
        const adios2::Box<adios2::Dims> sel({{Nx * rank}, {Nx}});
        bpFloats.SetSelection(sel);
        bpReader.Get(bpFloats, simData.data());
    }
    auto bpStep = bpIO.InquireVariable<unsigned int>("bpStep");
    if (bpStep)
    {
        bpReader.Get(bpStep, &inStep);
    }

    bpReader.EndStep();
}

```

17. Now we will close the file.

```
bpReader.Close();
```

18. Now we will call the writer and reader functions:

```

writer(adios, engine, filename, Nx, nSteps, rank, size);
reader(adios, engine, filename, Nx, nSteps, rank, size);

```

19. Finally we need to finalize MPI.

```
MPI_Finalize();
```

20. The final code should look as follows (excluding try/catch and optional usage of MPI), and it was derived from the example `ADIOS2/examples/hello/bpStepsWriteRead/bpStepsWriteRead.cpp`.

```

/*
 * Distributed under the OSI-approved Apache License, Version 2.0. See
 * accompanying file Copyright.txt for details.
 *
 * bpStepsWriteRead.cpp Simple example of writing and reading data through ADIOS2 BP_
 * engine with
 * multiple simulations steps for every IO step.
 *
 * Created on: Feb 16, 2017
 * Author: William F Godoy godoywf@ornl.gov
 */

#include <algorithm> // std::for_each
#include <ios>       // std::ios_base::failure
#include <iostream>  // std::cout
#if ADIOS2_USE_MPI
#include <mpi.h>
#endif
#include <stdexcept> //std::invalid_argument std::exception

```

(continues on next page)

(continued from previous page)

```

#include <vector>

#include <adios2.h>

void update_array(std::vector<float> &array, int val)
{
    std::transform(array.begin(), array.end(), array.begin(),
        [val](float v) -> float { return v + static_cast<float>(val); });
}

void writer(adios2::ADIOS &adios, const std::string &engine, const std::string &fname,
    const size_t Nx, unsigned int nSteps, int rank, int size)
{
    std::vector<float> simData(Nx, 0);

    adios2::IO bpIO = adios.DeclareIO("WriteIO");
    bpIO.SetEngine(engine);

    const adios2::Dims shape{static_cast<size_t>(size * Nx)};
    const adios2::Dims start{static_cast<size_t>(rank * Nx)};
    const adios2::Dims count{Nx};
    auto bpFloats = bpIO.DefineVariable<float>("bpFloats", shape, start, count);

    auto bpStep = bpIO.DefineVariable<unsigned int>("bpStep");

    adios2::Engine bpWriter = bpIO.Open(fname, adios2::Mode::Write);

    for (unsigned int step = 0; step < nSteps; ++step)
    {
        const adios2::Box<adios2::Dims> sel({0}, {Nx});
        bpFloats.SetSelection(sel);

        bpWriter.BeginStep();
        bpWriter.Put(bpFloats, simData.data());
        bpWriter.Put(bpStep, step);
        bpWriter.EndStep();

        // Update values in the simulation data
        update_array(simData, 10);
    }

    bpWriter.Close();
}

void reader(adios2::ADIOS &adios, const std::string &engine, const std::string &fname,
    const size_t Nx, unsigned int /*nSteps*/, int rank, int /*size*/)
{
    adios2::IO bpIO = adios.DeclareIO("ReadIO");
    bpIO.SetEngine(engine);

    adios2::Engine bpReader = bpIO.Open(fname, adios2::Mode::Read);

```

(continues on next page)

(continued from previous page)

```

std::vector<float> simData(Nx, 0);
unsigned int inStep = 0;
for (unsigned int step = 0; bpReader.BeginStep() == adios2::StepStatus::OK; ++step)
{
    auto bpFloats = bpIO.InquireVariable<float>("bpFloats");
    if (bpFloats)
    {
        const adios2::Box<adios2::Dims> sel({{Nx * rank}, {Nx}});
        bpFloats.SetSelection(sel);
        bpReader.Get(bpFloats, simData.data());
    }
    auto bpStep = bpIO.InquireVariable<unsigned int>("bpStep");
    if (bpStep)
    {
        bpReader.Get(bpStep, &inStep);
    }

    bpReader.EndStep();
    if (inStep != step)
    {
        std::cout << "ERROR: step mismatch\n";
        return;
    }
}
bpReader.Close();
}

int main(int argc, char *argv[])
{
    int rank, size;

#ifdef ADIOS2_USE_MPI
    int provided;
    // MPI_THREAD_MULTIPLE is only required if you enable the SST MPI_DP
    MPI_Init_thread(&argc, &argv, MPI_THREAD_MULTIPLE, &provided);
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
#else
    rank = 0;
    size = 1;
#endif

    const std::string engine = argv[1] ? argv[1] : "BPFile";
    std::cout << "Using engine " << engine << std::endl;

    const std::string filename = engine + "StepsWriteRead.bp";
    const unsigned int nSteps = 10;
    const unsigned int Nx = 60000;
    try
    {
        /** ADIOS class factory of IO class objects */
#ifdef ADIOS2_USE_MPI

```

(continues on next page)



(continued from previous page)

```

        adios2::ADIOS adios(MPI_COMM_WORLD);
#else
        adios2::ADIOS adios;
#endif

        writer(adios, engine, filename, Nx, nSteps, rank, size);
        reader(adios, engine, filename, Nx, nSteps, rank, size);
    }
    catch (std::invalid_argument &e)
    {
        std::cout << "Invalid argument exception, STOPPING PROGRAM from rank " << rank <
        ↪ << "\n";
        std::cout << e.what() << "\n";
    }
    catch (std::ios_base::failure &e)
    {
        std::cout << "IO System base failure exception, STOPPING PROGRAM "
                    << "from rank "
                    << rank << "\n";
        std::cout << e.what() << "\n";
    }
    catch (std::exception &e)
    {
        std::cout << "Exception, STOPPING PROGRAM from rank " << rank << "\n";
        std::cout << e.what() << "\n";
    }

    #if ADIOS2_USE_MPI
        MPI_Finalize();
    #endif

    return 0;
}

```

21. You can compile and run it as follows:

```

cd Path-To-ADIOS2/examples/hello/bpStepsWriteRead
mkdir build
cd build
cmake -DADIOS2_DIR=Path-To-ADIOS2/build/ ..
cmake --build .
mpirun -np 2 ./adios2_hello_bpStepsWriteRead_mpi

```

22. You can check the content of the output file “BPFileStepsWriteRead.bp” using *bpls* as follows:

```

Path-To-ADIOS2/build/bin/bpls ./BPFileStepsWriteRead.bp

float      bpFloats  10*{1200000}
uint32_t   bpStep    10*scalar

```



## HDF5 API SUPPORT THROUGH VOL

We have developed a HDF5 VOL in order to comply with the ECP request to support HDF5 API. Through this VOL the HDF5 clients can read and write general ADIOS files.

### 25.1 Disclaimer

---

**Note:** The Virtual Object Layer (VOL) is a feature introduced in recent release of HDF5 1.12 ([https://hdf5.wiki/index.php/New\\_Features\\_in\\_HDF5\\_Release\\_1.12](https://hdf5.wiki/index.php/New_Features_in_HDF5_Release_1.12)).

So please do make sure your HDF5 version supports the latest VOL.

---

Once the ADIOS VOL is compiled, There are two ways to apply it:

- externally (through dynamic library, no code change)
- internally (through client code).

### 25.2 External

- Set the following environment parameters:

```
HDF5_VOL_CONNECTOR=ADIOS2_VOL  
HDF5_PLUGIN_PATH=/replace/with/your/adios2_vol/lib/path/
```

Without code change, ADIOS2 VOL will be loaded at runtime by HDF5, to access ADIOS files without changing user code.

### 25.3 Internal

- include adios header
- call the function to set VOL when H5F is initiated
- call the function to unset VOL when H5F is closed

```
// other includes
#include <adios2/h5vol/H5Vol.h> // ADD THIS LINE TO INCLUDE ADIOS VOL

hid_t pid = H5Pcreate(H5P_FILE_ACCESS);
// other declarations
hid_t fid = H5Fopen(filename, mode, pid);

H5VL_ADIOS2_set(pid); // ADD THIS LINE TO USE ADIOS VOL

H5Fclose(fid);

H5VL_ADIOS2_unset(); // ADD THIS LINE TO EXIT ADIOS VOL
```

**Note:** The following features are not supported in this VOL:

- hyperslab support
- HDF5 parameters are not passed down. e.g. compression/decompression
- ADIOS2 parameters is not setup.
- user defined types
- change of variable extent is not supported in ADIOS2.

## COMMAND LINE UTILITIES

ADIOS 2 provides a set of command line utilities for quick data exploration and manipulation that builds on top of the library. They are located inside the `adios2-install-location/bin` directory after a `make install`.

**Tip:** Optionally the `adios2-install-location/bin` location can be added to your `PATH` to avoid absolute paths when using `adios2` command line utilities.

Currently supported tools are:

- `bppls` : exploration of `bp/hdf5` files data and metadata in human readable formats
- `adios_reorganize`
- `adios2-config`
- `sst_conn_tool` : SST staging engine connectivity diagnostic tool

### 26.1 bppls : Inspecting Data

The `bppls` utility is for examining and pretty-printing the content of ADIOS output files (BP and HDF5 files). By default, it lists the variables in the file including the type, name, and dimensionality.

Let's assume we run the Heat Transfer tutorial example and produce the output by

```
$ mpirun -n 12 ./heatSimulation sim.bp 3 4 5 4 3 1
Process decomposition : 3 x 4
Array size per process : 5 x 4
Number of output steps : 3
Iterations per step : 1

$ mpirun -n 3 ./heatAnalysis sim.bp a.bp 3 1

$ bppls a.bp
double    T      3*{15, 16}
double    dT     3*{15, 16}
```

In our example, we have two arrays, `T` and `dT`. Both are 2-dimensional double arrays, their global size is 15x16 and the file contains 3 output steps of these arrays.

**Note:** `bppls` is written in C++ and therefore sees the order of the dimensions in *row major*. If the data was written from Fortran in column-major order, you will see the dimension order flipped when listing with `bppls`, just as a code written

in C++ or python would see the data.

---

Here is the description of the most used options (use `bpls -h` to print help on all options for this utility).

- `-l`

Print the min/max of the arrays and the values of scalar values

```
$ bpls -l a.bp
double   T           3*{15, 16} = 0 / 200
double   dT          3*{15, 16} = -53.1922 / 49.7861
```

- `-a -A`

List the attributes along with the variables. `-A` will print the attributes only.

```
$ bpls a.bp -la
double   T           3*{15, 16} = 0 / 200
string   T/description attr = "Temperature from simulation"
string   T/unit       attr = "C"
double   dT          3*{15, 16} = -53.1922 / 49.7861
string   dT/description attr = "Temperature difference between two steps,
↳calculated in analysis"
```

- `pattern, -e`

Select which variables/attributes to list or dump. By default the pattern(s) are interpreted as shell file patterns.

```
$ bpls a.bp -la T*
double   T           3*{15, 16} = 0 / 200
```

Multiple patterns can be defined in the command line.

```
$ bpls a.bp -la T/* dT/*
string   T/description attr = "Temperature from simulation"
string   T/unit       attr = "C"
string   dT/description attr = "Temperature difference between two steps,
↳calculated in analysis"
```

If the `-e` option is given (all) the pattern(s) will be interpreted as regular expressions.

```
$ bpls a.bp -la T.* -e
double   T           3*{15, 16} = 0 / 200
string   T/description attr = "Temperature from simulation"
string   T/unit       attr = "C"
```

- `-D`

Print the decomposition of a variable. In the BP file, the data blocks written by different writers are stored separately and have their own size info and min/max statistics. This option is useful at code development to check if the output file is written the way intended.

```
$ bpls a.bp -l T -D
double   T           3*{15, 16} = 0 / 200
step 0:
  block 0: [ 0: 4, 0:15] = 3.54199e-14 / 200
```

(continues on next page)

(continued from previous page)

```

block 1: [ 5: 9, 0:15] = 58.3642 / 200
block 2: [10:14, 0:15] = 0 / 200
step 1:
block 0: [ 0: 4, 0:15] = 31.4891 / 153.432
block 1: [ 5: 9, 0:15] = 68.2107 / 180.184
block 2: [10:14, 0:15] = 31.4891 / 161.699
step 2:
block 0: [ 0: 4, 0:15] = 48.0431 / 135.225
block 1: [ 5: 9, 0:15] = 74.064 / 170.002
block 2: [10:14, 0:15] = 48.0431 / 147.87

```

In this case we find 3 blocks per output step and 3 output steps. We can see that the variable T was decomposed in the first (slow) dimension. In the above example, the T variable in the simulation output (sim.bp) had 12 blocks per step, but the analysis code was running on 3 processes, effectively reorganizing the data into fewer larger blocks.

- -d

Dump the data content of a variable. For pretty-printing, one should use the additional -n and -f options. For selecting only a subset of a variable, one should use the -s and -c options.

By default, six values are printed per line and using C style -g prints for floating point values.

```

$ bpls a.bp -d T
double T      3*{15, 16}
(0, 0, 0)     124.925 124.296 139.024 95.2078 144.864 191.485
(0, 0, 6)     139.024 140.814 124.925 109.035 110.825 58.3642
(0, 0, 12)    104.985 154.641 110.825 125.553 66.5603 65.9316
...
(2, 14, 4)    105.918 116.842 111.249 102.044 93.3121 84.5802
(2, 14, 10)   75.3746 69.782 80.706 93.5492 94.7595 95.0709

```

For pretty-printing, use the additional -n and -f options.

```

$ bpls a.bp -d T -n 16 -f "%3.0f"
double T      3*{15, 16}
(0, 0, 0)     125 124 139 95 145 191 139 141 125 109 111 58 105 155 111 126
(0, 1, 0)      67 66 81 37 86 133 81 82 67 51 52 0 47 96 52 67
(0, 2, 0)     133 133 148 104 153 200 148 149 133 118 119 67 114 163 119 134
...
(2, 13, 0)      98 98 96 96 115 132 124 109 97 86 71 63 79 98 97 95
(2, 14, 0)      96 96 93 93 106 117 111 102 93 85 75 70 81 94 95 95

```

For selecting a subset of a variable, use the -s and -c options. These options are N+1 dimensional for N-dimensional arrays with more than one steps. The first element of the options are used to select the starting step and the number of steps to print.

The following example dumps a 4x4 small subset from the center of the array, one step from the second (middle) step:

```

$ bpls a.bp -d T -s "1,6,7" -c "1,4,4" -n 4
double T      3*{15, 16}
slice (1:1, 6:9, 7:10)
(1,6, 7)      144.09 131.737 119.383 106.787
(1,7, 7)      145.794 133.44 121.086 108.49

```

(continues on next page)

(continued from previous page)

```
(1,8, 7)    145.794 133.44 121.086 108.49
(1,9, 7)    144.09 131.737 119.383 106.787
```

- `-y --noindex`

Data can be dumped in a format that is easier to import later into other tools, like Excel. The leading array indexes can be omitted by using this option. Non-data lines, like the variable and slice info, are printed with a starting `;`.

```
$ bpls a.bp -d T -s "1,6,7" -c "1,4,4" -n 4 --noindex
; double   T      3*{15, 16}
; slice (1:1, 6:9, 7:10)
144.09 131.737 119.383 106.787
145.794 133.44 121.086 108.49
145.794 133.44 121.086 108.49
144.09 131.737 119.383 106.787
```

**Note:** HDF5 files can also be dumped with bpls if ADIOS was built with HDF5 support. Note that the HDF5 files do not contain min/max information for the arrays and therefore bpls always prints 0 for them:

```
$ bpls -l a.h5
double   T      3*{15, 16} = 0 / 0
double   dT     3*{15, 16} = 0 / 0
```

## 26.2 adios\_reorganize

`adios_reorganize` and `adios_reorganize_mpi` are generic ADIOS tools to read in ADIOS streams and output the same data into another ADIOS stream. The two tools are for serial and MPI environments, respectively. They can be used for

- converting files between ADIOS BP and HDF5 format
- using separate compute nodes to stage I/O from/to disk to/from a large scale application
- reorganizing the data blocks for a different number of blocks

Let's assume we run the Heat Transfer tutorial example and produce the output by

```
$ mpirun -n 12 ./heatSimulation sim.bp 3 4 5 4 3 1
Process decomposition : 3 x 4
Array size per process : 5 x 4
Number of output steps : 3
Iterations per step : 1

$ bpls sim.bp
double   T      3*{15, 16}
```

In our example, we have an array, `T`. It is a 2-dimensional double array, its global size is 15x16 and the file contains 3 output steps of this array. The array is composed of 12 separate blocks coming from the 12 producers in the application.

- Convert BP file to HDF5 file



If ADIOS is built with HDF5 support, this tool can be used to convert between the two file formats.

```
$ mpirun -n 2 adios_reorganize_mpi sim.bp sim.h5 BPFFile "" HDF5 "" 2 1

$ bpls sim.h5
double T 3*{15, 16}

$ h5ls -r sim.h5
/
/Step0
/Step0/T
/Step1
/Step1/T
/Step2
/Step2/T
Group
Group
Dataset {15, 16}
Group
Dataset {15, 16}
Group
Dataset {15, 16}
```

- Stage I/O through extra compute nodes

If writing data to disk is a bottleneck to the application, it may be worth to use extra nodes for receiving the data quickly from the application and then write to disk while the application continues computing. Similarly, data can be staged in from disk into extra nodes and make it available for fast read-in for an application. One can use one of the staging engines in ADIOS to perform this data staging (SST, SSC, DataMan).

Assuming that the heatSimulation is using SST instead of file I/O in a run (set in its `adios2.xml` configuration file), staging to disk can be done this way:

```
Make sure adios2.xml sets SST for the simulation:
<io name="SimulationOutput">
  <engine type="SST">
  </engine>
</io>

$ mpirun -n 12 ./heatSimulation sim.bp 3 4 5 4 3 1 : \
        -n 2 adios_reorganize_mpi sim.bp staged.bp SST "" BPFFile "" 2 1

$ bpls staged.bp
double T 3*{15, 16}
```

Data is staged to the extra 2 cores and those will write the data to disk while the heatSimulation calculates the next step. Note, that this staging can only be useful if the tool can write all data to disk before the application produces the next output step. Otherwise, it will still block the application for I/O.

- Reorganizing the data blocks in file for a different number of blocks

In the above example, the application writes the array from 12 processes, but then `adios_reorganize_mpi` reads the global arrays on 2 processes. The output file on disk will therefore contain the array in 2 blocks. This reorganization of the array may be useful if reading is too slow for a dataset created by many-many processes. One may want to reorganize a file written by tens or hundreds of thousands of processes if one wants to read the content more than one time and the read time proves to be a bottleneck in one's work flow.

```
$ mpirun -n 12 ./heatSimulation sim.bp 3 4 5 4 3 1
$ bpls sim.bp -D
```

(continues on next page)

(continued from previous page)

```

double T      3*{15, 16}
  step 0:
    block 0: [ 0: 4,  0: 3]
    block 1: [ 5: 9,  0: 3]
    block 2: [10:14,  0: 3]
    block 3: [ 0: 4,  4: 7]
    block 4: [ 5: 9,  4: 7]
    block 5: [10:14,  4: 7]
    block 6: [ 0: 4,  8:11]
    block 7: [ 5: 9,  8:11]
    block 8: [10:14,  8:11]
    block 9: [ 0: 4, 12:15]
    block 10: [ 5: 9, 12:15]
    block 11: [10:14, 12:15]
  step 1:
    block 0: [ 0: 4,  0: 3]
    block 1: [ 5: 9,  0: 3]
    block 2: [10:14,  0: 3]
    block 3: [ 0: 4,  4: 7]
    block 4: [ 5: 9,  4: 7]
    block 5: [10:14,  4: 7]
    block 6: [ 0: 4,  8:11]
    block 7: [ 5: 9,  8:11]
    block 8: [10:14,  8:11]
    block 9: [ 0: 4, 12:15]
    block 10: [ 5: 9, 12:15]
    block 11: [10:14, 12:15]
  step 2:
    block 0: [ 0: 4,  0: 3]
    block 1: [ 5: 9,  0: 3]
    block 2: [10:14,  0: 3]
    block 3: [ 0: 4,  4: 7]
    block 4: [ 5: 9,  4: 7]
    block 5: [10:14,  4: 7]
    block 6: [ 0: 4,  8:11]
    block 7: [ 5: 9,  8:11]
    block 8: [10:14,  8:11]
    block 9: [ 0: 4, 12:15]
    block 10: [ 5: 9, 12:15]
    block 11: [10:14, 12:15]

$ mpirun -n 2 adios_reorganize_mpi sim.bp reorg.bp BPFile "" BPFile "" 2 1
$ bpls reorg.bp -D
double T      3*{15, 16}
  step 0:
    block 0: [ 0: 6,  0:15]
    block 1: [ 7:14,  0:15]
  step 1:
    block 0: [ 0: 6,  0:15]
    block 1: [ 7:14,  0:15]
  step 2:

```

(continues on next page)

(continued from previous page)

```
block 0: [ 0: 6,  0:15]
block 1: [ 7:14, 0:15]
```

## 26.3 adios2-config

*adios2-config* is provided to aid with non-CMake builds (e.g. manually generated Makefile). Running the *adios2-config* command under *adios2-install-dir/bin/adios2-config* will generate the following usage information:

```
./adios2-config --help
adios2-config (--help | [--c-flags] [--c-libs] [--cxx-flags] [--cxx-libs] [--fortran-
  flags] [--fortran-libs])
--help          Display help information
-c             Both compile and link flags for the C bindings
--c-flags       Preprocessor and compile flags for the C bindings
--c-libs        Linker flags for the C bindings
-x            Both compile and link flags for the C++ bindings
--cxx-flags     Preprocessor and compile flags for the C++ bindings
--cxx-libs      Linker flags for the C++ bindings
-f            Both compile and link flags for the F90 bindings
--fortran-flags Preprocessor and compile flags for the F90 bindings
--fortran-libs  Linker flags for the F90 bindings
```

Please refer to the *From non-CMake build systems* for more information on how to use this command.

## 26.4 sst\_conn\_tool : SST network connectivity tool

The *sst\_conn\_tool* utility exposes some aspects of SST network connectivity parameters and activity in order to allow debugging of SST connections.

In its simplest usage, it just lets you test an SST connection (between two runs of the program) and tells you the network information its trying, I.E. what IP address and port it determined to use for listening, and if it's connecting somewhere what those parameters are. For example, you'd first run *sst\_conn\_tool* in one window and its output would look like this:

```
bash-3.2$ bin/sst_conn_tool

Sst writer is listening for TCP/IP connection at IP 192.168.1.17, port 26051

Sst connection tool waiting for connection...
```

To try to connect from another window, you run *sst\_conn\_tool* with the *-c* or *—connect* option:

```
bash-3.2$ bin/sst_conn_tool -c

Sst reader at IP 192.168.1.17, listening UDP port 26050

Attempting TCP/IP connection to writer at IP 192.168.1.17, port 26051
```

(continues on next page)

(continued from previous page)

```
Connection success, all is well!
bash-3.2$
```

Here, it has found the contact information file, tried and succeeded in making the connection and has indicated that all is well. In the first window, we get a similar message about the success of the connection.

In the event that there is trouble with the connection, there is a “-i” or “-info” option that will provide additional information about the network configuration options. For example:

```
bash-3.2$ bin/sst_conn_tool -i

ADIOS2_IP_CONFIG best guess hostname is "sandy.local"
ADIOS2_IP_CONFIG Possible interface lo0 : IPV4 127.0.0.1
ADIOS2_IP_CONFIG Possible interface en0 : IPV4 192.168.1.56
ADIOS2_IP_CONFIG Possible interface en5 : IPV4 192.168.1.17
ADIOS2_IP_CONFIG best guess IP is "192.168.1.17"
ADIOS2_IP_CONFIG default port range is "any"

The following environment variables can impact ADIOS2_IP_CONFIG operation:
    ADIOS2_IP           - Publish the specified IP address for contact
    ADIOS2_HOSTNAME     - Publish the specified hostname for contact
    ADIOS2_USE_HOSTNAME - Publish a hostname preferentially over IP_
↪address
    ADIOS2_INTERFACE   - Use the IP address associated with the_
↪specified network interface
    ADIOS2_PORT_RANGE  - Use a port within the specified range
↪"low:high",
                        or specify "any" to let the OS choose

Sst writer is listening for TCP/IP connection at IP 192.168.1.17, port 26048

Sst connection tool waiting for connection...
```

Full options for sst\_conn\_tool:

Operational Modes:

- **-l --listen**  
Display connection parameters and wait for an SST connection (default)
- **-c --connect** Attempt a connection to an already-running instance of sst\_conn\_tool

Additional Options:

- **-i --info**  
Display additional networking information for this host
- **-f --file**  
Use file-based contact info sharing (default). The SST contact file is created in the current directory
- **-s --screen**  
Use screen-based contact info sharing, SST contact info is displayed/entered via terminal
- **-h --help**  
Display this message usage and options

## VISUALIZING DATA

Certain ADIOS2 bp files can be recognized by third party visualization tools. This section describes how to create an ADIOS2 bp file to accomodate the visualization product requirements.

### 27.1 Using VTK and Paraview

ADIOS BP files can now be seamlessly integrated into the [Visualization Toolkit](#) and [Paraview](#). Datasets can be described with an additional attribute that conforms to the [VTK XML data model formats](#) as high-level descriptors that will allow interpretation of ADIOS2 variables into a hierarchy understood by the VTK infrastructure. This XML data format is saved in ADIOS2 as either an attribute or as an additional `vtk.xml` file inside the `file.bp.dir` directory.

There are currently a number of limitations:

- It only works with MPI builds of VTK and Paraview
- Support only one block per ADIOS dataset
- Only supports BP Files, streams are not supported
- Currently working up to 3D (and linearized 1D) variables for scalars and vectors.
- Image Data, vti, is supported with ADIOS2 Global Array Variables only
- Unstructured Grid, vtu, is supported with ADIOS2 Local Arrays Variables only

Two VTK file types are supported:

1. Image data (.vti)
2. Unstructured Grid (.vtu)

The main idea is to populate the above XML format contents describing the extent and the data arrays with ADIOS variables in the BP data set. The result is a more-than-compact typical VTK data file since extra information about the variable, such as dimensions and type, as they already exist in the BP data set.

A typical VTK image data XML descriptor (.vti):

```
<?xml version="1.0"?>
<VTKFile type="ImageData">
  <ImageData WholeExtent="x1 x2 y1 y2 z1 z2" Origin="x0 y0 z0" Spacing="dx dy dz">
    <Piece Extent="x1 x2 y1 y2 z1 z2">
      <PointData>
        <DataArray Name="p1"/>
        <DataArray Name="p2"/>
      </PointData>
      <CellData>
```

(continues on next page)

(continued from previous page)

```

    <DataArray Name="c1"/>
    <DataArray Name="c2"/>
  </CellData>
</Piece>
</ImageData>
</VTKFile>

```

A typical VTK unstructured grid XML descriptor (.vtu):

```

<?xml version="1.0"?>
<VTKFile type="ImageData">
  <ImageData WholeExtent="x1 x2 y1 y2 z1 z2" Origin="x0 y0 z0" Spacing="dx dy dz">
    <Piece Extent="x1 x2 y1 y2 z1 z2">
      <PointData>
        <DataArray Name="p1"/>
        <DataArray Name="p2"/>
      </PointData>
      <CellData>
        <DataArray Name="c1"/>
        <DataArray Name="c2"/>
      </CellData>
    </Piece>
  </ImageData>
</VTKFile>

```

In addition, VTK can interpret physical-time or output-step varying data stored with ADIOS by reusing the special “TIME” tag. This is better illustrated in the following section.

### 27.1.1 Saving the vtk.xml data model

For the full source code of the following illustration example see the [gray-scott adios2 tutorial](#).

To incorporate the data model in a BP data file, the application has two options:

- 1) Adding a string attribute called “vtk.xml” in code. “TIME” is a special tag for adding physical time variables.

```

const std::string imageData = R"(
  <?xml version="1.0"?>
  <VTKFile type="ImageData" version="0.1" byte_order="LittleEndian">
    <ImageData WholeExtent=")" + extent + R"(" Origin="0 0 0" Spacing="1 1 1">
      <Piece Extent=")" + extent + R"(">
        <CellData Scalars="U">
          <DataArray Name="U" />
          <DataArray Name="V" />
          <DataArray Name="TIME">
            step
          </DataArray>
        </CellData>
      </Piece>
    </ImageData>
  </VTKFile>)" );

```

```
io.DefineAttribute<std::string>("vtk.xml", imageData);
```

**Tip:** C++11 users should take advantage C++11 string literals (`R"( xml_here )"`) to simplify escaping characters in the XML.

The resulting bpls output should contain the “vtk.xml” attribute and the variables in the model:

```
> bpls gs.bp -lav
File info:
  of variables: 3
  of attributes: 7
  statistics:   Min / Max

double  Du      attr  = 0.2
double  Dv      attr  = 0.1
double  F       attr  = 0.02
double  U       24*{48, 48, 48} = 0.107439 / 1.04324
double  V       24*{48, 48, 48} = 0 / 0.672232
double  dt      attr  = 1
double  k       attr  = 0.048
double  noise   attr  = 0.01
int32_t step    24*scalar = 0 / 575
string  vtk.xml attr  =

<VTKFile type="ImageData" version="0.1" byte_order="LittleEndian">
  <ImageData WholeExtent="0 49 0 49 0 49" Origin="0 0 0" Spacing="1 1 1">
    <Piece Extent="0 49 0 49 0 49">
      <CellData Scalars="U">
        <DataArray Name="U" />
        <DataArray Name="V" />
        <DataArray Name="TIME">
          step
        </DataArray>
      </CellData>
    </Piece>
  </ImageData>
</VTKFile>
```

2) Saving a “vtk.xml” file inside the file.bp.dir to describe the data after is created

```
> cat gs.bp.dir/vtk.xml

<?xml version="1.0"?>
<VTKFile type="ImageData" version="0.1" byte_order="LittleEndian">
  <ImageData WholeExtent=")" + extent + R"( " Origin="0 0 0" Spacing="1 1 1">
    <Piece Extent=")" + extent + R"(">
      <CellData Scalars="U">
        <DataArray Name="U" />
        <DataArray Name="V" />
        <DataArray Name="TIME">
          step
        </DataArray>
      </CellData>
    </Piece>
  </ImageData>
```

(continues on next page)

(continued from previous page)

&lt;/VTKFile&gt;

This BP file should be recognized by Paraview:

Similarly, unstructured grid (.vtu) support can be added with the limitations of using specific labels for the variable names setting the “connectivity”, “vertices”, and cell “types”.

The following example is taken from example 2 of the [MFEM product examples website](#) using ADIOS2:

The resulting *bpls* output for unstructured grid data types:

```
File info:
  of variables: 6
  of attributes: 4
  statistics:   Min / Max

uint32_t  NumOfElements      {4} = 1024 / 1024
uint32_t  NumOfVertices      {4} = 1377 / 1377
string    app                attr  = "MFEM"
uint64_t  connectivity       [4]*{1024, 9} = 0 / 1376
uint32_t  dimension          attr  = 3
string    glvis_mesh_version attr  = "1.0"
double    sol                 [4]*{1377, 3} = -0.201717 / 1.19304
uint32_t  types              scalar = 11
double    vertices           [4]*{1377, 3} = -1.19304 / 8.20172
string    vtk.xml            attr  =

<VTKFile type="UnstructuredGrid" version="0.1" byte_order="LittleEndian">
  <UnstructuredGrid>
    <Piece NumberOfPoints="NumOfVertices" NumberOfCells="NumOfElements">
      <Points>
        <DataArray Name="vertices" />
      </Points>
      <Cells>
        <DataArray Name="connectivity" />
        <DataArray Name="types" />
      </Cells>
      <PointData>
        <DataArray Name="sol" />
      </PointData>
    </Piece>
  </UnstructuredGrid>
</VTKFile>
```

and resulting visualization in Paraview for different “cell” types:



## 28.1 MPI vs Non-MPI

1. *Can I use the same library for MPI and non-MPI code?*

## 28.2 APIs

1. *Can I use ADIOS 2 C++11 library with C++98 codes?*
2. *Why are C and Fortran APIs missing functionality?*
3. *C++11: Why are `std::string` arguments passed sometimes by value and sometimes by reference?*
4. *C++11: Should I pass `adios2::` objects by value or by reference?*
5. *Fortran: Can I pass slices and temporary arrays to `adios2_put`?*

## 28.3 Building on Titan

1. *My application uses PGI compilers on Titan, can I link ADIOS 2?*
2. *How do I enable the Python bindings on Titan?*

## 28.4 Building and Running on Fujitsu FX100

1. *How do I build ADIOS 2 on Fujitsu FX100?*
2. *SST engine hangs on Fujitsu FX100. Why?*

## 28.5 FAQs Answered

### 28.5.1 Can I use the same library for MPI and non-MPI code?

Short answer: Yes, since version 2.6.0.

Long answer: One build of ADIOS2 can be used by both serial and parallel applications. Use the `-s` and `-m` flags in the `adios2-config` command. By default, or with the `-m` flag, the command gives the flags for a parallel build, which add `-DADIOS2_USE_MPI` to the compilation flags and include extra libraries containing the MPI implementations into the linker flags. The `-s` flag will omit these flags. For example, if ADIOS2 is installed into `/opt/adios2`, the flags for a Fortran application will look like these:

```
$ /opt/adios2/bin/adios2-config --fortran-flags
-DADIOS2_USE_MPI -I/opt/adios2/include/adios2/fortran
$ /opt/adios2/bin/adios2-config --fortran-flags -m
-DADIOS2_USE_MPI -I/opt/adios2/include/adios2/fortran
$ /opt/adios2/bin/adios2-config --fortran-flags -s
-I/opt/adios2/include/adios2/fortran

$ /opt/adios2/bin/adios2-config --fortran-libs
-Wl,-rpath,/opt/adios2/lib /opt/adios2/lib/libadios2_fortran_mpi.so.2.6.0 /opt/
↪adios2/lib/libadios2_fortran.so.2.6.0 -Wl,-rpath-link,/opt/adios2/lib
$ /opt/adios2/bin/adios2-config --fortran-libs -s
-Wl,-rpath,/opt/adios2/lib /opt/adios2/lib/libadios2_fortran.so.2.6.0 -Wl,-
↪rpath-link,/opt/adios2/lib
```

If using `cmake`, there are different targets to build parallel

```
find_package(MPI REQUIRED)
find_package(ADIOS2 REQUIRED)
#...
add_library(my_library src1.cxx src2.cxx)
target_link_libraries(my_library PRIVATE adios2::cxx11_mpi MPI::MPI_C)
#...
add_library(my_f_library src1.F90 src2.F90)
target_link_libraries(my_f_library PRIVATE adios2::fortran_mpi adios2::fortran_
↪MPI::MPI_Fortran)
```

and serial applications:

```
find_package(ADIOS2 REQUIRED)
#...
add_library(my_library src1.cxx src2.cxx)
target_link_libraries(my_library PRIVATE adios2::cxx11)
#...
add_library(my_f_library src1.F90 src2.F90)
target_link_libraries(my_f_library PRIVATE adios2::fortran)
```

### 28.5.2 Can I use ADIOS 2 C++11 library with C++98 codes?

Use the *C bindings*. C++11 is a brand new language standard and many new (and old, *e.g.* `std::string`) might cause ABI conflicts.

### 28.5.3 Why are C and Fortran APIs missing functionality?

Because language intrinsics are NOT THE SAME. For example, C++ and Python support key/value pair structures natively, *e.g.* `std::map` and dictionaries, respectively. Fortran and C only support arrays natively. Use the right language (tool) for the right task.

### 28.5.4 C++11: Why are `std::string` arguments passed sometimes by value and sometimes by reference?

C++11, provides mechanisms to optimize copying small objects, rather than passing by reference. The latter was always the rule for C++98. When a string is passed by value, it's assumed that the name will be short,  $\leq 15$  characters, most of the time. While passing by reference indicates that the string can be of any size. Check the [isocpp guidelines on this topic](#) for more information.

### 28.5.5 C++11: Should I pass `adios2::` objects by value or by reference?

`adios2::ADIOS`: always pass by reference this is the only “large memory” object; all others: pass by reference or value depending on your coding standards and requirements, they are small objects that wrap around a pointer to an internal object inside `adios2::ADIOS`.

### 28.5.6 Fortran: Can I pass slices and temporary arrays to `adios2_put`?

By definition the lifetime of a temporary is the scope of the function is passed to. Therefore, you must use sync mode with `adios2_put`. Deferred mode will save garbage data since the memory location of a temporary is undefined after `adios2_put`, not able to reach `adios2_end_step`, `adios2_close` or `adios2_perform_puts` where the memory is actually used.

### 28.5.7 My application uses PGI compilers on Titan, can I link ADIOS 2?

Follow directions at [Building on HPC Systems](#) to setup support for PGI on Titan. PGI compilers depend on GNU headers, but they must point to a version greater than gcc 4.8.1 to support C++11 features. The gcc module doesn't need to be loaded, though. Example:

```
$ module load gcc/7.2.0
$ makelocalrc $(dirname $(which pgc++)) -gcc $(which gcc) -gpp $(which g++) -
→g77 $(which gfortran) -o -net 1>${HOME}/.mypgirc 2>/dev/null
$ module unload gcc/7.2.0
```

## 28.5.8 How do I enable the Python bindings on Titan?

The default ADIOS2 configuration on Titan builds a static library. Python bindings require enabling the dynamic libraries and the Cray dynamic environment variable. See *Building on HPC Systems* and *Enabling the Python bindings*. For example:

```
[atkins3@titan-ext4 code]$ mkdir adios
[atkins3@titan-ext4 code]$ cd adios
[atkins3@titan-ext4 adios]$ git clone https://github.com/ornladios/adios2.git
→source
[atkins3@titan-ext4 adios]$ module swap PrgEnv-pgi PrgEnv-gnu
[atkins3@titan-ext4 adios]$ module load cmake3/3.11.3
[atkins3@titan-ext4 adios]$ module load python python_numpy python_mpi4py
[atkins3@titan-ext4 adios]$ export CRAYPE_LINK_TYPE=dynamic CC=cc CXX=CC FC=ftn
[atkins3@titan-ext4 adios]$ mkdir build
[atkins3@titan-ext4 build]$ cd build
[atkins3@titan-ext4 build]$ cmake ../source
-- The C compiler identification is GNU 6.3.0
-- The CXX compiler identification is GNU 6.3.0
-- Cray Programming Environment 2.5.13 C
-- Check for working C compiler: /opt/cray/craype/2.5.13/bin/cc
-- Check for working C compiler: /opt/cray/craype/2.5.13/bin/cc -- works
-- Detecting C compiler ABI info
-- Detecting C compiler ABI info - done
-- Detecting C compile features
-- Detecting C compile features - done
-- Cray Programming Environment 2.5.13 CXX
-- Check for working CXX compiler: /opt/cray/craype/2.5.13/bin/CC
-- Check for working CXX compiler: /opt/cray/craype/2.5.13/bin/CC -- works
...
-- Found PythonInterp: /sw/titan/.swci/0-login/opt/spack/20180315/linux-suse_
→linux11-x86_64/gcc-4.3.4/python-2.7.9-v6ctjewwdx6k2qs7ublexz7gnx457jo5/bin/
→python2.7 (found version "2.7.9")
-- Found PythonLibs: /sw/titan/.swci/0-login/opt/spack/20180315/linux-suse_
→linux11-x86_64/gcc-4.3.4/python-2.7.9-v6ctjewwdx6k2qs7ublexz7gnx457jo5/lib/
→libpython2.7.so (found version "2.7.9")
-- Found PythonModule_numpy: /sw/xk6/python_numpy/1.7.1/python2.7.9_craylibsci_
→gnu4.9.0/lib64/python2.7/site-packages/numpy
-- Found PythonModule_mpi4py: /lustre/atlas/sw/xk7/python_mpi4py/2.0.0/cle5.
→2up04_python2.7.9/lib64/python2.7/site-packages/mpi4py
-- Found PythonFull: /sw/titan/.swci/0-login/opt/spack/20180315/linux-suse_
→linux11-x86_64/gcc-4.3.4/python-2.7.9-v6ctjewwdx6k2qs7ublexz7gnx457jo5/bin/
→python2.7 found components: Interp Libs numpy mpi4py
...
ADIOS2 build configuration:
  ADIOS Version: 2.4.0
  C++ Compiler : GNU 6.3.0 CrayPrgEnv
    /opt/cray/craype/2.5.13/bin/CC

  Fortran Compiler : GNU 6.3.0 CrayPrgEnv
    /opt/cray/craype/2.5.13/bin/ftn

  Installation prefix: /usr/local
```

(continues on next page)

(continued from previous page)

```

    bin: bin
    lib: lib
include: include
    cmake: lib/cmake/adios2
    python: lib/python2.7/site-packages

Features:
  Library Type: shared
  Build Type:   Release
  Testing: ON
  Build Options:
    BZip2      : ON
    ZFP        : OFF
    SZ         : OFF
    MGARD      : OFF
    MPI        : ON
    DataMan    : ON
    SST        : ON
    ZeroMQ     : OFF
    HDF5       : OFF
    Python     : ON
    Fortran    : ON
    SysVShMem  : ON
    Endian_Reverse: OFF

-- Configuring done
-- Generating done
-- Build files have been written to: /ccs/home/atkins3/code/adios/build

```

### 28.5.9 How do I build ADIOS 2 on Fujitsu FX100?

- Cross-compilation (building on the login node) is not recommended. Submit an interactive job and build on the compute nodes.
- Make sure CMake >= 3.6 is installed on the compute nodes. If not, you need to build and install it from source since CMake does not provide SPARC V9 binaries.
- Use gcc instead of the Fujitsu compiler. We tested with gcc 6.3.0
- CMake fails to automatically find the correct MPI library on FX100. As a workaround, set CC, CXX, and FC to the corresponding MPI compiler wrappers:

```
$ CC=mpigcc CXX=mpig++ FC=mpigfortran cmake ..
```

### 28.5.10 SST engine hangs on Fujitsu FX100. Why?

The communication thread of SST might have failed to start. FX100 requires users to set the maximum stack size manually when launching POSIX threads. One way to do this is through `ulimit` (*e.g.* `ulimit -s 1024`). You can also set the stack size when submitting the job. Please contact your system administrator for details.

## ADVICE

This list is similar to the Advice sections for each chapter in [The C++ Programming Language, Fourth Edition](#) by Bjarne Stroustrup. The goal is to provide specific advice and good practices about the use of ADIOS2 in other projects.

1. Use `MPI_COMM_SELF` to run MPI compiled versions of ADIOS 2 in “serial” mode
2. Use a runtime configuration file in the ADIOS constructor or `adios2_init` when targeting multiple engines
3. Check object validity when developing (similar to `fstream`):

- C++: *operator bool*

```
if(var) engine.Put(var, data);
```

- C: NULL pointer

```
if(var) adios2_put(engine, var, data, adios2_mode_deferred);
```

- Python: `v2 __nonzero__ v3 __bool__`. Note: do not check for None object

```
if(var) engine.Put(var, data);
```

- Fortran: `type%valid`

```
if( adios%valid .eqv. .true. ) then  
    adios2_declare_io(adios, io, "foo")  
end if
```

4. C++11 and Python: use try-catch (try-except in Python) blocks to handle exceptions from ADIOS 2
5. C++11: use fixed-width types (`int8_t`, `uint32_t`) for portability
6. Define your data structure: set of variables and attributes before developing. Data hierarchies/models can be built on top of ADIOS 2.
7. Read the documentation for [Supported Engines](#) before targeting development for a particular engine
8. MPI development: treat ADIOS constructor/destructor (`adios2_init/adios2_finalize`) and Engine Open and Close always as collective functions. For the most part, ADIOS 2 API functionality is local, but other Engine functions might follow other rules, [Supported Engines](#).
9. Use *Remove* functions carefully. They create dangling objects/pointers.
10. Thread-safety: treat ADIOS 2 as NOT thread-safe. Either use a mutex or only handle I/O from a master thread. ADIOS 2 is about performance, adding I/O serial algorithm operations into a parallel execution block may reduce parallel portions from Amdahl’s Law.

11. Prefer the high-level Python and C++ APIs for simple tasks that do not require performance. The more familiar Write/Read overloads for File I/O return native data constructs (`std::vector` and `numpy` arrays) immediately for a requested selection. `open` only explores the metadata index.
12. C++: prefer templates to `void*` to increase compile-time safety. Use `IO::InquireVariableType("variableName")` and `adios2::GetType<T>()` to cast upfront to a `Variable<T>`. C++17 has `std::any` as an alternative to `void*`. ADIOS 2 follows closely the STL model.
13. Understand Put and Get memory contracts from *Engine*
14. Prefer Put/Get Deferred mode, treat Sync as a special mode
15. Put Span: create all spans in a step before populating them. Spans follow the same iterator invalidation rules as `std::vector`, so use `span.data()` to always keep the span pointer up-to-date
16. Always populate data before calling Put in deferred mode, and do not change it between Put and EndStep, or Close
17. Never call `PerformPuts` right before `EndStep`. This was a code pattern that had no adverse effects with the BP3/4 file engines and is present in some older code, but was never beneficial.
18. Use `BeginStep` and `EndStep` to write code that is portable across all ADIOS 2 Engine types: file and streaming.
19. Always use `Close` for every call to `Open`.
20. C, Fortran: always call `adios2_finalize` for every call to `adios2_init` to avoid memory leaks.
21. Reminder: C++, C, Python: Row-Major, while Fortran: Column-Major. ADIOS 2 will handle interoperability between ordering. Remember that *bpls : Inspecting Data* is always a Row-Major reader so Fortran reader need to swap dimensions seen in bpls. bpls: (slow, ..., fast) -> Fortran(fast,...,slow).
22. Fortran API: use the type members (`var%valid`, `var%name`, etc.) to get extra type information.
23. Fortran C interoperability: Fortran bindings support the majority of applications using Fortran 90. We currently don't support the ISO\_C\_BINDING interoperability module in Fortran 2003.
24. Always keep the IO object self-contained keeping its own set of `Variables`, `Attributes` and `Engines`. Do not combine `Variables` with multiple `Engines` or multiple modes, unless it's 100% guaranteed to be safe in your program avoiding `Variable` access conflicts.
25. Developers: explore the testing infrastructure `ADIOS2/testing` in ADIOS 2 as a starting point for using ADIOS 2 in your own testing environment.
26. Become a super-user of *bpls : Inspecting Data* to analyze datasets generated by ADIOS 2.
  - search



## A

- `add_operation()` (*adios2.Variable method*), 181
- `add_operation_string()` (*adios2.Variable method*), 181
- `add_transport()` (*adios2.IO method*), 176
- `AddOperation()` (*adios2.bindings.Variable method*), 186
- `AddTransport()` (*adios2.bindings.IO method*), 185
- `adios` (*adios2.Stream property*), 171
- `Adios` (*class in adios2*), 175
- `ADIOS` (*class in adios2.bindings*), 184
- `adios()` (*adios2.IO method*), 176
- `adios2::ADIOS` (C++ class), 80
- `adios2::ADIOS::~~ADIOS` (C++ function), 81
- `adios2::ADIOS::ADIOS` (C++ function), 80, 81
- `adios2::ADIOS::AtIO` (C++ function), 81
- `adios2::ADIOS::DeclareIO` (C++ function), 81
- `adios2::ADIOS::DefineOperator` (C++ function), 82
- `adios2::ADIOS::EnterComputationBlock` (C++ function), 83
- `adios2::ADIOS::ExitComputationBlock` (C++ function), 83
- `adios2::ADIOS::FlushAll` (C++ function), 82
- `adios2::ADIOS::InquireOperator` (C++ function), 82
- `adios2::ADIOS::operator bool` (C++ function), 81
- `adios2::ADIOS::operator=` (C++ function), 81
- `adios2::ADIOS::RemoveAllIOs` (C++ function), 83
- `adios2::ADIOS::RemoveIO` (C++ function), 82
- `adios2::Attribute` (C++ class), 93
- `adios2::Attribute::Attribute` (C++ function), 93
- `adios2::Attribute::Data` (C++ function), 93
- `adios2::Attribute::IsValue` (C++ function), 93
- `adios2::Attribute::Name` (C++ function), 93
- `adios2::Attribute::operator bool` (C++ function), 93
- `adios2::Attribute::Type` (C++ function), 93
- `adios2::Engine` (C++ class), 94
- `adios2::Engine::~~Engine` (C++ function), 94
- `adios2::Engine::AllStepsBlocksInfo` (C++ function), 99
- `adios2::Engine::BeginStep` (C++ function), 94
- `adios2::Engine::BetweenStepPairs` (C++ function), 99
- `adios2::Engine::BlocksInfo` (C++ function), 100
- `adios2::Engine::Close` (C++ function), 99
- `adios2::Engine::CurrentStep` (C++ function), 95
- `adios2::Engine::EndStep` (C++ function), 99
- `adios2::Engine::Engine` (C++ function), 94
- `adios2::Engine::Flush` (C++ function), 99
- `adios2::Engine::Get` (C++ function), 96–99
- `adios2::Engine::GetAbsoluteSteps` (C++ function), 100
- `adios2::Engine::LockReaderSelections` (C++ function), 100
- `adios2::Engine::LockWriterDefinitions` (C++ function), 100
- `adios2::Engine::Name` (C++ function), 94
- `adios2::Engine::OpenMode` (C++ function), 94
- `adios2::Engine::operator bool` (C++ function), 94
- `adios2::Engine::PerformDataWrite` (C++ function), 96
- `adios2::Engine::PerformGets` (C++ function), 99
- `adios2::Engine::PerformPuts` (C++ function), 96
- `adios2::Engine::Put` (C++ function), 95, 96
- `adios2::Engine::Steps` (C++ function), 100
- `adios2::Engine::Type` (C++ function), 94
- `adios2::fstream` (C++ class), 155
- `adios2::fstream::~~fstream` (C++ function), 157
- `adios2::fstream::close` (C++ function), 164
- `adios2::fstream::current_step` (C++ function), 164
- `adios2::fstream::end_step` (C++ function), 164
- `adios2::fstream::fstream` (C++ function), 156, 157
- `adios2::fstream::getstep` (C++ function), 164
- `adios2::fstream::open` (C++ function), 157, 158
- `adios2::fstream::openmode` (C++ enum), 155
- `adios2::fstream::openmode::app` (C++ enumerator), 156
- `adios2::fstream::openmode::in` (C++ enumerator), 155
- `adios2::fstream::openmode::in_flatten_steps` (C++ enumerator), 156
- `adios2::fstream::openmode::in_random_access`

(C++ *enumerator*), 156  
adios2::fstream::openmode::out (C++ *enumerator*), 155  
adios2::fstream::operator bool (C++ *function*), 157  
adios2::fstream::read (C++ *function*), 160–163  
adios2::fstream::read\_attribute (C++ *function*), 163  
adios2::fstream::set\_parameter (C++ *function*), 158  
adios2::fstream::set\_parameters (C++ *function*), 158  
adios2::fstream::write (C++ *function*), 159, 160  
adios2::fstream::write\_attribute (C++ *function*), 158, 159  
adios2::Group (C++ *class*), 101  
adios2::Group::AttributeType (C++ *function*), 103  
adios2::Group::AvailableAttributes (C++ *function*), 102  
adios2::Group::AvailableGroups (C++ *function*), 102  
adios2::Group::AvailableVariables (C++ *function*), 102  
adios2::Group::InquireAttribute (C++ *function*), 102  
adios2::Group::InquireGroup (C++ *function*), 102  
adios2::Group::InquirePath (C++ *function*), 102  
adios2::Group::InquireVariable (C++ *function*), 102  
adios2::Group::setPath (C++ *function*), 102  
adios2::Group::VariableType (C++ *function*), 102  
adios2::IO (C++ *class*), 83  
adios2::IO::~~IO (C++ *function*), 83  
adios2::IO::AddOperation (C++ *function*), 88  
adios2::IO::AddTransport (C++ *function*), 84  
adios2::IO::AttributeType (C++ *function*), 88  
adios2::IO::AvailableAttributes (C++ *function*), 88  
adios2::IO::AvailableVariables (C++ *function*), 87  
adios2::IO::ClearParameters (C++ *function*), 84  
adios2::IO::DefineAttribute (C++ *function*), 85  
adios2::IO::DefineVariable (C++ *function*), 84  
adios2::IO::EngineType (C++ *function*), 88  
adios2::IO::FlushAll (C++ *function*), 87  
adios2::IO::InConfigFile (C++ *function*), 83  
adios2::IO::InquireAttribute (C++ *function*), 86  
adios2::IO::InquireGroup (C++ *function*), 87  
adios2::IO::InquireVariable (C++ *function*), 85  
adios2::IO::IO (C++ *function*), 83  
adios2::IO::Name (C++ *function*), 83  
adios2::IO::Open (C++ *function*), 87  
adios2::IO::operator bool (C++ *function*), 83  
adios2::IO::Parameters (C++ *function*), 84  
adios2::IO::RemoveAllAttributes (C++ *function*), 87  
adios2::IO::RemoveAllVariables (C++ *function*), 86  
adios2::IO::RemoveAttribute (C++ *function*), 86  
adios2::IO::RemoveVariable (C++ *function*), 86  
adios2::IO::SetEngine (C++ *function*), 83  
adios2::IO::SetParameter (C++ *function*), 83  
adios2::IO::SetParameters (C++ *function*), 84  
adios2::IO::SetTransportParameter (C++ *function*), 84  
adios2::IO::VariableType (C++ *function*), 88  
adios2::Operator (C++ *class*), 100  
adios2::Operator::Operator (C++ *function*), 101  
adios2::Operator::operator bool (C++ *function*), 101  
adios2::Operator::Parameters (C++ *function*), 101  
adios2::Operator::SetParameter (C++ *function*), 101  
adios2::Operator::Type (C++ *function*), 101  
adios2::Variable (C++ *class*), 89  
adios2::Variable::~~Variable (C++ *function*), 89  
adios2::Variable::AddOperation (C++ *function*), 91  
adios2::Variable::AllStepsBlocksInfo (C++ *function*), 92  
adios2::Variable::BlockID (C++ *function*), 91  
adios2::Variable::Count (C++ *function*), 90  
adios2::Variable::GetAccuracy (C++ *function*), 92  
adios2::Variable::GetMemorySpace (C++ *function*), 89  
adios2::Variable::Info (C++ *struct*), 92  
adios2::Variable::Info::BlockID (C++ *member*), 93  
adios2::Variable::Info::Count (C++ *member*), 92  
adios2::Variable::Info::Data (C++ *function*), 92  
adios2::Variable::Info::IsReverseDims (C++ *member*), 93  
adios2::Variable::Info::IsValue (C++ *member*), 93  
adios2::Variable::Info::Max (C++ *member*), 92  
adios2::Variable::Info::Min (C++ *member*), 92  
adios2::Variable::Info::Start (C++ *member*), 92  
adios2::Variable::Info::Step (C++ *member*), 93  
adios2::Variable::Info::Value (C++ *member*), 92  
adios2::Variable::Info::WriterID (C++ *member*), 93  
adios2::Variable::Max (C++ *function*), 92  
adios2::Variable::Min (C++ *function*), 91  
adios2::Variable::MinMax (C++ *function*), 91  
adios2::Variable::Name (C++ *function*), 90  
adios2::Variable::Operations (C++ *function*), 91  
adios2::Variable::operator bool (C++ *function*), 89

adios2::Variable::RemoveOperations (C++ function), 91  
 adios2::Variable::SelectionSize (C++ function), 90  
 adios2::Variable::SetAccuracy (C++ function), 90  
 adios2::Variable::SetBlockSelection (C++ function), 89  
 adios2::Variable::SetMemorySelection (C++ function), 89  
 adios2::Variable::SetMemorySpace (C++ function), 89  
 adios2::Variable::SetSelection (C++ function), 89  
 adios2::Variable::SetShape (C++ function), 89  
 adios2::Variable::SetStepSelection (C++ function), 90  
 adios2::Variable::Shape (C++ function), 90  
 adios2::Variable::ShapeID (C++ function), 90  
 adios2::Variable::Sizeof (C++ function), 90  
 adios2::Variable::Start (C++ function), 90  
 adios2::Variable::Steps (C++ function), 91  
 adios2::Variable::StepsStart (C++ function), 91  
 adios2::Variable::Type (C++ function), 90  
 adios2::Variable::Variable (C++ function), 89  
 adios2\_add\_operation (C++ function), 143  
 adios2\_add\_transport (C++ function), 134  
 adios2\_at\_io (C++ function), 131  
 adios2\_attribute\_data (C++ function), 146  
 adios2\_attribute\_is\_value (C++ function), 145  
 adios2\_attribute\_name (C++ function), 145  
 adios2\_attribute\_size (C++ function), 145  
 adios2\_attribute\_type (C++ function), 145  
 adios2\_attribute\_type\_string (C++ function), 145  
 adios2\_available\_attributes (C++ function), 138  
 adios2\_available\_variables (C++ function), 137  
 adios2\_begin\_step (C++ function), 147  
 adios2\_between\_step\_pairs (C++ function), 147  
 adios2\_clear\_parameters (C++ function), 134  
 adios2\_close (C++ function), 150  
 adios2\_close\_by\_index (C++ function), 150  
 adios2\_current\_step (C++ function), 147  
 adios2\_declare\_io (C++ function), 131  
 adios2\_declare\_io\_order (C++ function), 131  
 adios2\_define\_attribute (C++ function), 135  
 adios2\_define\_attribute\_array (C++ function), 135  
 adios2\_define\_operator (C++ function), 131  
 adios2\_define\_variable (C++ function), 134  
 adios2\_define\_variable\_attribute (C++ function), 136  
 adios2\_define\_variable\_attribute\_array (C++ function), 136  
 adios2\_end\_step (C++ function), 149  
 adios2\_engine\_get\_type (C++ function), 146  
 adios2\_engine\_name (C++ function), 146  
 adios2\_engine\_openmode (C++ function), 146  
 adios2\_engine\_type (C++ function), 139  
 adios2\_enter\_computation\_block (C++ function), 132  
 adios2\_exit\_computation\_block (C++ function), 132  
 adios2\_finalize (C++ function), 132  
 adios2\_flush (C++ function), 149  
 adios2\_flush\_all (C++ function), 131  
 adios2\_flush\_all\_engines (C++ function), 139  
 adios2\_flush\_by\_index (C++ function), 149  
 adios2\_free\_blockinfo (C++ function), 150  
 adios2\_get (C++ function), 148  
 adios2\_get\_by\_name (C++ function), 149  
 adios2\_get\_engine (C++ function), 139  
 adios2\_get\_memory\_space (C++ function), 140  
 adios2\_get\_parameter (C++ function), 133  
 adios2\_in\_config\_file (C++ function), 133  
 adios2\_init (C macro), 130  
 adios2\_init\_config (C macro), 130  
 adios2\_init\_config\_mpi (C++ function), 130  
 adios2\_init\_config\_serial (C++ function), 130  
 adios2\_init\_mpi (C++ function), 130  
 adios2\_init\_serial (C++ function), 130  
 adios2\_inquire\_all\_attributes (C++ function), 137  
 adios2\_inquire\_all\_variables (C++ function), 135  
 adios2\_inquire\_attribute (C++ function), 136  
 adios2\_inquire\_blockinfo (C++ function), 150  
 adios2\_inquire\_group\_attributes (C++ function), 137  
 adios2\_inquire\_group\_variables (C++ function), 135  
 adios2\_inquire\_operator (C++ function), 131  
 adios2\_inquire\_subgroups (C++ function), 137  
 adios2\_inquire\_variable (C++ function), 135  
 adios2\_inquire\_variable\_attribute (C++ function), 137  
 adios2\_lock\_reader\_selections (C++ function), 150  
 adios2\_lock\_writer\_definitions (C++ function), 150  
 adios2\_open (C++ function), 138  
 adios2\_open\_new\_comm (C++ function), 138  
 adios2\_operator\_type (C++ function), 151  
 adios2\_perform\_data\_write (C++ function), 148  
 adios2\_perform\_gets (C++ function), 149  
 adios2\_perform\_puts (C++ function), 148  
 adios2\_put (C++ function), 147  
 adios2\_put\_by\_name (C++ function), 148  
 adios2\_remove\_all\_attributes (C++ function), 138  
 adios2\_remove\_all\_ios (C++ function), 132  
 adios2\_remove\_all\_variables (C++ function), 137

`adios2_remove_attribute` (C++ function), 138  
`adios2_remove_io` (C++ function), 132  
`adios2_remove_operations` (C++ function), 144  
`adios2_remove_variable` (C++ function), 137  
`adios2_selection_size` (C++ function), 143  
`adios2_set_block_selection` (C++ function), 140  
`adios2_set_engine` (C++ function), 133  
`adios2_set_memory_selection` (C++ function), 140  
`adios2_set_memory_space` (C++ function), 140  
`adios2_set_operation_parameter` (C++ function), 144  
`adios2_set_parameter` (C++ function), 133  
`adios2_set_parameters` (C++ function), 133  
`adios2_set_selection` (C++ function), 140  
`adios2_set_shape` (C++ function), 139  
`adios2_set_step_selection` (C++ function), 141  
`adios2_set_transport_parameter` (C++ function), 134  
`adios2_steps` (C++ function), 147  
`adios2_variable_count` (C++ function), 143  
`adios2_variable_max` (C++ function), 144  
`adios2_variable_min` (C++ function), 144  
`adios2_variable_name` (C++ function), 141  
`adios2_variable_ndims` (C++ function), 142  
`adios2_variable_shape` (C++ function), 142  
`adios2_variable_shape_with_memory_space` (C++ function), 142  
`adios2_variable_shapeid` (C++ function), 142  
`adios2_variable_start` (C++ function), 142  
`adios2_variable_steps` (C++ function), 143  
`adios2_variable_steps_start` (C++ function), 143  
`adios2_variable_type` (C++ function), 141  
`adios2_variable_type_string` (C++ function), 141  
`all_blocks_info` (adios2.Engine method), 179  
`all_blocks_info` (adios2.Stream method), 171  
`at_io` (adios2.Adios method), 175  
`AtIO` (adios2.bindings.ADIOS method), 184  
`Attribute` (class in adios2), 183  
`Attribute` (class in adios2.bindings), 187  
`available_attributes` (adios2.IO method), 176  
`available_attributes` (adios2.Stream method), 171  
`available_variables` (adios2.IO method), 176  
`available_variables` (adios2.Stream method), 171  
`AvailableAttributes` (adios2.bindings.IO method), 185  
`AvailableVariables` (adios2.bindings.IO method), 185

## B

`begin_step` (adios2.Engine method), 179  
`begin_step` (adios2.Stream method), 172  
`BeginStep` (adios2.bindings.Engine method), 187  
`between_step_pairs` (adios2.Engine method), 179

`BetweenStepPairs` (adios2.bindings.Engine method), 187  
`block_id` (adios2.Variable method), 181  
`BlockID` (adios2.bindings.Variable method), 186  
`blocks_info` (adios2.Engine method), 179  
`BlocksInfo` (adios2.bindings.Engine method), 187

## C

`Close` (adios2.bindings.Engine method), 187  
`close` (adios2.Engine method), 179  
`close` (adios2.Stream method), 172  
`Count` (adios2.bindings.Variable method), 186  
`count` (adios2.Variable method), 181  
`current_step` (adios2.Engine method), 179  
`current_step` (adios2.Stream method), 172  
`CurrentStep` (adios2.bindings.Engine method), 187

## D

`data` (adios2.Attribute method), 183  
`Data` (adios2.bindings.Attribute method), 187  
`data_string` (adios2.Attribute method), 183  
`DataStream` (adios2.bindings.Attribute method), 187  
`declare_io` (adios2.Adios method), 175  
`DeclareIO` (adios2.bindings.ADIOS method), 184  
`define_attribute` (adios2.IO method), 177  
`define_operator` (adios2.Adios method), 175  
`define_variable` (adios2.IO method), 177  
`define_variable` (adios2.Stream method), 172  
`DefineAttribute` (adios2.bindings.IO method), 185  
`DefineOperator` (adios2.bindings.ADIOS method), 184  
`DefineVariable` (adios2.bindings.IO method), 185

## E

`end_step` (adios2.Engine method), 180  
`end_step` (adios2.Stream method), 172  
`EndStep` (adios2.bindings.Engine method), 187  
`engine` (adios2.Stream property), 172  
`Engine` (class in adios2), 179  
`Engine` (class in adios2.bindings), 187  
`engine_type` (adios2.IO method), 177  
`EngineType` (adios2.bindings.IO method), 185

## F

`FileReader` (class in adios2), 175  
`Flush` (adios2.bindings.Engine method), 187  
`flush` (adios2.Engine method), 180  
`flush_all` (adios2.Adios method), 175  
`flush_all` (adios2.IO method), 177  
`FlushAll` (adios2.bindings.ADIOS method), 184  
`FlushAll` (adios2.bindings.IO method), 185

## G

`Get` (adios2.bindings.Engine method), 187



get() (*adios2.Engine method*), 180  
 get\_parameters() (*adios2.Operator method*), 183  
 GetBlockIDs() (*adios2.bindings.Query method*), 188  
 GetResult() (*adios2.bindings.Query method*), 188

## I

impl (*adios2.Adios property*), 175  
 impl (*adios2.Attribute property*), 183  
 impl (*adios2.Engine property*), 180  
 impl (*adios2.IO property*), 178  
 impl (*adios2.Operator property*), 183  
 impl (*adios2.Variable property*), 181  
 inquire\_attribute() (*adios2.IO method*), 178  
 inquire\_attribute() (*adios2.Stream method*), 172  
 inquire\_operator() (*adios2.Adios method*), 175  
 inquire\_variable() (*adios2.IO method*), 178  
 inquire\_variable() (*adios2.Stream method*), 172  
 InquireAttribute() (*adios2.bindings.IO method*), 185  
 InquireOperator() (*adios2.bindings.ADIOS method*), 184  
 InquireVariable() (*adios2.bindings.IO method*), 185  
 io (*adios2.Stream property*), 173  
 IO (*class in adios2*), 176  
 IO (*class in adios2.bindings*), 185

## L

lock\_reader\_selections() (*adios2.Engine method*), 180  
 lock\_writer\_definitions() (*adios2.Engine method*), 180  
 LockReaderSelections() (*adios2.bindings.Engine method*), 187  
 LockWriterDefinitions() (*adios2.bindings.Engine method*), 187  
 loop\_index() (*adios2.Stream method*), 173

## M

mode (*adios2.Stream property*), 173

## N

name() (*adios2.Attribute method*), 183  
 Name() (*adios2.bindings.Attribute method*), 187  
 Name() (*adios2.bindings.Engine method*), 187  
 Name() (*adios2.bindings.Variable method*), 186  
 name() (*adios2.Variable method*), 181  
 num\_steps() (*adios2.Stream method*), 173

## O

Open() (*adios2.bindings.IO method*), 185  
 open() (*adios2.IO method*), 178  
 Operations() (*adios2.bindings.Variable method*), 186  
 operations() (*adios2.Variable method*), 181

Operator (*class in adios2*), 183  
 Operator (*class in adios2.bindings*), 188

## P

Parameters() (*adios2.bindings.IO method*), 185  
 Parameters() (*adios2.bindings.Operator method*), 188  
 parameters() (*adios2.IO method*), 178  
 perform\_data\_write() (*adios2.Engine method*), 180  
 perform\_gets() (*adios2.Engine method*), 180  
 perform\_puts() (*adios2.Engine method*), 180  
 PerformDataWrite() (*adios2.bindings.Engine method*), 187  
 PerformGets() (*adios2.bindings.Engine method*), 187  
 PerformPuts() (*adios2.bindings.Engine method*), 187  
 Put() (*adios2.bindings.Engine method*), 188  
 put() (*adios2.Engine method*), 180

## Q

Query (*class in adios2.bindings*), 188

## R

read() (*adios2.Stream method*), 173  
 read\_attribute() (*adios2.Stream method*), 173  
 read\_attribute\_string() (*adios2.Stream method*), 174  
 remove\_all\_attributes() (*adios2.IO method*), 178  
 remove\_all\_ios() (*adios2.Adios method*), 176  
 remove\_all\_variables() (*adios2.IO method*), 178  
 remove\_attribute() (*adios2.IO method*), 178  
 remove\_io() (*adios2.Adios method*), 176  
 remove\_operations() (*adios2.Variable method*), 181  
 remove\_variable() (*adios2.IO method*), 179  
 RemoveAllAttributes() (*adios2.bindings.IO method*), 185  
 RemoveAllIOs() (*adios2.bindings.ADIOS method*), 184  
 RemoveAllVariables() (*adios2.bindings.IO method*), 185  
 RemoveAttribute() (*adios2.bindings.IO method*), 186  
 RemoveIO() (*adios2.bindings.ADIOS method*), 184  
 RemoveOperations() (*adios2.bindings.Variable method*), 186  
 RemoveVariable() (*adios2.bindings.IO method*), 186

## S

selection\_size() (*adios2.Variable method*), 181  
 SelectionSize() (*adios2.bindings.Variable method*), 186  
 set\_block\_selection() (*adios2.Variable method*), 181  
 set\_engine() (*adios2.IO method*), 179  
 set\_parameter() (*adios2.IO method*), 179  
 set\_parameter() (*adios2.Operator method*), 183  
 set\_parameters() (*adios2.IO method*), 179

`set_parameters()` (*adios2.Stream method*), 174  
`set_selection()` (*adios2.Variable method*), 181  
`set_shape()` (*adios2.Variable method*), 182  
`set_step_selection()` (*adios2.Variable method*), 182  
`SetBlockSelection()` (*adios2.bindings.Variable method*), 186  
`SetEngine()` (*adios2.bindings.IO method*), 186  
`SetParameter()` (*adios2.bindings.IO method*), 186  
`SetParameter()` (*adios2.bindings.Operator method*), 188  
`SetParameters()` (*adios2.bindings.IO method*), 186  
`SetSelection()` (*adios2.bindings.Variable method*), 186  
`SetShape()` (*adios2.bindings.Variable method*), 186  
`SetStepSelection()` (*adios2.bindings.Variable method*), 186  
`Shape()` (*adios2.bindings.Variable method*), 186  
`shape()` (*adios2.Variable method*), 182  
`shape_id()` (*adios2.Variable method*), 182  
`ShapeID()` (*adios2.bindings.Variable method*), 186  
`single_value()` (*adios2.Attribute method*), 183  
`single_value()` (*adios2.Variable method*), 182  
`SingleValue()` (*adios2.bindings.Attribute method*), 187  
`SingleValue()` (*adios2.bindings.Variable method*), 186  
`Sizeof()` (*adios2.bindings.Variable method*), 186  
`sizeof()` (*adios2.Variable method*), 182  
`Start()` (*adios2.bindings.Variable method*), 186  
`start()` (*adios2.Variable method*), 182  
`step_status()` (*adios2.Stream method*), 174  
`Steps()` (*adios2.bindings.Engine method*), 188  
`Steps()` (*adios2.bindings.Variable method*), 186  
`steps()` (*adios2.Engine method*), 180  
`steps()` (*adios2.Stream method*), 174  
`steps()` (*adios2.Variable method*), 182  
`steps_start()` (*adios2.Variable method*), 182  
`StepsStart()` (*adios2.bindings.Variable method*), 186  
`Stream` (*class in adios2*), 171

## T

`type()` (*adios2.Attribute method*), 183  
`Type()` (*adios2.bindings.Attribute method*), 187  
`Type()` (*adios2.bindings.Engine method*), 188  
`Type()` (*adios2.bindings.Operator method*), 188  
`Type()` (*adios2.bindings.Variable method*), 186  
`type()` (*adios2.Variable method*), 182

## V

`Variable` (*class in adios2*), 180  
`Variable` (*class in adios2.bindings*), 186

## W

`write()` (*adios2.Stream method*), 174  
`write_attribute()` (*adios2.Stream method*), 175