



# CGNS Proposal for Extension 0050

## Solution Continuity Type

Version 1.0

April 22, 2026

<b>CPEX#</b>	0050
<b>Scope</b>	Distinguish continuous and discontinuous solution continuity
<b>Champion</b>	M. Scot Breitenfeld
<b>Organization</b>	The HDF Group
<b>E-mail</b>	brtnfld@hdfgroup.org
<b>GitHub Issue</b>	TBD
<b>Reference Impl.</b>	TBD
<b>Target Release</b>	TBD
<b>Date First Posted</b>	April 22, 2026
<b>SIDS Status</b>	proposed
<b>Filemap Status</b>	proposed
<b>MLL Status</b>	proposed

## Contents

<b>1 Abstract</b>	<b>3</b>	9
<b>2 Motivation and Rationale</b>	<b>3</b>	10
2.1 Problem Statement . . . . .	3	11
2.2 Inadequacy of Existing CGNS Constructs . . . . .	4	12
2.3 Generality Beyond Higher-Order Methods . . . . .	4	13
<b>3 Detailed Description</b>	<b>4</b>	14
3.1 Design Overview . . . . .	4	15
3.2 New Enumeration: <code>SolutionContinuity_t</code> . . . . .	5	16
3.2.1 Rationale for Two-Value Classification . . . . .	5	17
3.3 SIDS Node Definition . . . . .	6	18
3.3.1 Node Placement Rationale . . . . .	7	19
3.4 New MLL Functions . . . . .	8	20
3.4.1 C API . . . . .	8	21
3.4.2 Fortran API . . . . .	9	22

3.5	Internal Data Structure Changes . . . . .	10	23
3.6	Internal Lookup Table Initialization . . . . .	10	24
3.7	MLL Implementation Requirements . . . . .	10	25
3.7.1	Internal Tree Parser Update ( <code>cgns_internals.c</code> ) . . . . .	10	26
3.7.2	Read Function Error Handling . . . . .	11	27
3.7.3	Write Function Implementation . . . . .	11	28
3.8	Data Layout for Discontinuous Solutions at Shared Mesh Vertices . . . . .	13	29
3.9	Interaction with <code>Rind_t</code> (Ghost Layers) . . . . .	14	30
<b>4</b>	<b>Interaction with CPEX 0045 (Higher-Order Elements)</b>	<b>14</b>	<b>31</b>
<b>5</b>	<b>Backward Compatibility</b>	<b>15</b>	<b>32</b>
<b>6</b>	<b>Examples</b>	<b>15</b>	<b>33</b>
6.1	Example 1: Writing a Continuous Galerkin Solution (C) . . . . .	15	34
6.2	Example 2: Writing a Discontinuous Galerkin Solution (C) . . . . .	15	35
6.3	Example 3: Reading and Dispatching on Continuity (C) . . . . .	16	36
6.4	Example 4: HDG Solver with Both Continuity Types (C) . . . . .	16	37
6.5	Example 5: Fortran Usage . . . . .	16	38
6.6	Example 6: Low-Order Finite Volume (DG $p = 0$ ) . . . . .	17	39
6.7	Example 7: DG Solution on a Shared-Vertex Mesh (C) . . . . .	17	40
6.8	Example 8: <code>DiscreteData_t</code> with Discontinuous Error Indicator (C) . . . . .	18	41
6.9	Example 9: <code>ZoneSubRegion_t</code> with Boundary Flux Data (C) . . . . .	18	42
6.10	Example 10: Solver Restart with Continuity-Aware Read (C) . . . . .	19	43
6.11	Example 11: FSI Coupling — Structural and Fluid Solutions (C) . . . . .	19	44
6.12	Example 12: Conjugate Heat Transfer — Perfect and Imperfect Thermal Contact (C) . . . . .	20	45
<b>7</b>	<b>SIDS Impact</b>	<b>21</b>	<b>46</b>
<b>8</b>	<b>File Mapping Impact</b>	<b>21</b>	<b>47</b>
<b>9</b>	<b>Design Decisions</b>	<b>22</b>	<b>48</b>
<b>10</b>	<b>Summary of API Additions</b>	<b>23</b>	<b>49</b>
<b>11</b>	<b>References</b>	<b>24</b>	<b>50</b>

## 1 Abstract

This CPEX introduces a `SolutionContinuity_t` enumeration and corresponding SIDS node to the CGNS standard, thereby enabling CGNS files to specify whether solution data exhibits continuous or discontinuous field continuity. This distinction, most commonly associated with Continuous Galerkin (CG) and Discontinuous Galerkin (DG) finite element methods, is essential for the correct interpretation of solution data, yet is not addressed by any existing CGNS construct. The proposal is broadly applicable: it is valid for any polynomial order, any discretization family (including finite element, finite volume, and spectral element methods), and integrates seamlessly with the higher-order solution infrastructure introduced by CPEX 0045.

## 2 Motivation and Rationale

### 2.1 Problem Statement

Modern CFD and computational mechanics codes increasingly use discretization methods whose solution data cannot be fully described by the existing CGNS metadata. The core issue is the treatment of degrees of freedom (DOFs) at shared geometric entities (vertices, edges, faces):

- In a **continuous** representation (e.g., Continuous Galerkin FEM,  $C^0$  spectral element methods), DOFs at shared interfaces are *shared*: a single value exists per geometric node, and the global solution field is at least  $C^0$ -continuous across element boundaries.
- In a **discontinuous** representation (e.g., Discontinuous Galerkin, finite volume,  $C^{-1}$  spectral element methods), DOFs are *element-local*: each element owns its own values at every point, including at shared interfaces. Multiple independent values may exist at the same geometric location.

This distinction affects:

1. **Data size and layout.** A DG solution on a mesh with shared vertices stores more values than a CG solution on the same mesh. Post-processors that assume shared-vertex semantics will misinterpret the data dimensions.
2. **Visualization and interpolation.** CG fields can be interpolated globally using the mesh connectivity; DG fields require element-local interpolation and may exhibit intentional discontinuities at element interfaces that should not be smoothed.
3. **Restart and data exchange.** Solvers reading restart files must know whether DOFs are shared or element-local to correctly reconstruct the solution state. Coupling frameworks exchanging solution data between codes need this metadata to perform correct inter-code transfers. In Fluid-Structure Interaction (FSI) applications, for example, a continuous FEM structural displacement field and a discontinuous DG fluid field coexist—the coupling framework must use different interpolation and projection operators for each, and the continuity metadata enables this dispatch.
4. **Error estimation and adaptation.** The jump in solution values across element interfaces in a DG solution is physically meaningful and is used for error estimation. A post-processor that incorrectly treats DG data as continuous would average out these jumps.

## 2.2 Inadequacy of Existing CGNS Constructs 89

No existing CGNS mechanism captures the continuous/discontinuous distinction: 90

`GridLocation_t` specifies *where* data is located (vertices, cell centers, faces, edges, or interpolation points), but not whether values at shared locations are shared or element-local. A `Vertex-`located solution could be either CG or DG; the same applies to `InterpolationPoints`. 91  
92  
93

`PointSet` (`PointList`, `PointRange`) can index subsets of data but does not encode continuity semantics. 94  
95

`InterpolationType_t` (CPEX 0045) describes the polynomial basis (e.g., Lagrange, monomial), but not whether that basis yields a globally continuous or element-local field. 96  
97

`SolutionInterpolation_t` (CPEX 0045) records polynomial order and basis type per element type but contains no continuity information. 98  
99

`Data array size` can sometimes be used to *infer* the representation—a DG field will have  $N_{\text{elem}} \times N_{\text{DOF/elem}}$  values versus  $N_{\text{nodes}}$  for CG. However, this inference is fragile, ambiguous for  $p = 0$  (cell-centered) data, and breaks entirely for mixed-element meshes or when point sets are used. 100  
101  
102

## 2.3 Generality Beyond Higher-Order Methods 103

Although often discussed in the context of high-order finite elements, the continuous/discontinuous distinction is **not limited to higher-order methods**: 104  
105

- **DG**  $p = 0$  (piecewise constant) is essentially a cell-centered finite volume representation. 106
- **DG**  $p = 1$  (piecewise linear, discontinuous) is widely used and is a low-order method. 107
- **CG**  $p = 1$  is standard continuous linear FEM, the most common finite element discretization. 108
- Classical **finite volume** representations are inherently discontinuous ( $p = 0$  DG). 109
- **Hybridizable DG (HDG)** methods have both element-local DOFs (discontinuous) and trace DOFs on faces (continuous on the skeleton). 110  
111
- **FSI and multi-physics coupling** often combines a continuous FEM structural solver (shared displacement DOFs) with a discontinuous DG or FV fluid solver (element-local flow variables) within the same file or across coupled zones. 112  
113  
114

Accordingly, a proper metadata field should therefore live at the `FlowSolution_t` level, applicable at any order, not solely within the higher-order interpolation infrastructure. 115  
116

## 3 Detailed Description 117

### 3.1 Design Overview 118

The proposal introduces: 119

1. A new `SolutionContinuity_t` enumeration with values covering the common continuity types. 120
2. A new optional child node of `FlowSolution_t`, `DiscreteData_t`, and `ZoneSubRegion_t` recording the solution continuity. 121  
122
3. Mid-Level Library functions to read and write the continuity. 123

## 4. Fortran bindings. 124

## 3.2 New Enumeration: SolutionContinuity\_t 125

```

1  typedef enum {
2  CGNS_ENUMV( SolutionContinuityNull )           = CG_Null,
3  CGNS_ENUMV( SolutionContinuityUserDefined )   = CG_UserDefined,
4  CGNS_ENUMV( SolutionContinuous )              = 2,
5  CGNS_ENUMV( SolutionDiscontinuous )          = 3
6  } CGNS_ENUMT( SolutionContinuity_t );
7
8  #define NofValidSolutionContinuity 4
9
10 extern CGNSDLL const char *
11     SolutionContinuityName[NofValidSolutionContinuity];
12
13 CGNSDLL const char *cg_SolutionContinuityName(
14     CGNS_ENUMT(SolutionContinuity_t) type);

```

The enumeration values are: 126

**SolutionContinuityNull** No continuity specified. This is the default for backward compatibility: existing files and codes that do not set this field are unaffected. 127 128

**SolutionContinuityUserDefined** Application-defined continuity not covered by the standard values. When this value is used, writers *should* attach a `Descriptor_t` child node named "ContinuityDescription" to the parent `FlowSolution_t`, `DiscreteData_t`, or `ZoneSubRegion_t` node. The string value should identify the continuity convention (e.g., "HDG-trace", "C1-spline"). This allows readers to recognize the convention without requiring a separate side-channel. 129 130 131 132 133 134

**SolutionContinuous** DOFs at shared mesh entities (vertices, edges, faces) are *shared* between adjacent elements: each shared entity carries exactly one value. This is a statement about *storage and DOF-sharing semantics*, not about the mathematical smoothness of the field. Consequently,  $\mathbf{H}(\text{div})$ -conforming spaces (e.g., Raviart–Thomas) and  $\mathbf{H}(\text{curl})$ -conforming spaces (e.g., Nédélec) should also use this label: their DOFs are shared at element interfaces (normal-flux or tangential-component DOFs, respectively), even though the full vector field is not pointwise  $C^0$ -continuous. Typical of: Continuous Galerkin FEM,  $C^0$  spectral element methods, isogeometric analysis with  $C^0$  junctions,  $\mathbf{H}(\text{div})$  and  $\mathbf{H}(\text{curl})$  conforming spaces. 135 136 137 138 139 140 141 142

**SolutionDiscontinuous** The solution field is element-local: each element owns independent DOFs, including at shared geometric locations, and no interface DOF is shared between adjacent elements. Typical of: Discontinuous Galerkin methods, finite volume methods,  $C^{-1}$  spectral element methods. 143 144 145 146

## 3.2.1 Rationale for Two-Value Classification 147

The binary continuous/discontinuous classification (enum values `SolutionContinuous` and `SolutionDiscontinuous`) is deliberate. Methods that appear to require finer distinctions are already handled by the existing two values, combined with standard CGNS patterns: 148 149 150

**HDG and mixed-continuity methods** use both element-local volume DOFs and globally continuous trace DOFs. These are stored as separate `FlowSolution_t` nodes—one marked `SolutionDiscontinuous`, the other `SolutionContinuous`—following the same pattern 151 152 153

CGNS already uses to separate fields at different `GridLocation_t` values. A dedicated `MixedContinuity` value would conflate two distinct data layouts into a single node, making it harder, not easier, for post-processors to interpret the data.

**Penalized / weakly-continuous DG methods** enforce continuity through penalty terms in the variational formulation, but the data layout remains element-local: each element owns its own DOFs. The correct label is `SolutionDiscontinuous`. “Weak continuity” is a solver-internal property of the numerical method, not a property of the stored data that readers and post-processors need to act on.

**Isogeometric analysis** ( $C^k$ ,  $k \geq 1$ ) uses basis functions with higher-than- $C^0$  smoothness, but the data layout is still globally continuous with shared DOFs at patch interfaces. The correct label is `SolutionContinuous`. The smoothness order  $k$  is a property of the basis, not of the data layout; it belongs in the interpolation metadata (CPEX 0045), not in the continuity type.

The `SolutionContinuityUserDefined` value remains available as an escape hatch, consistent with CGNS convention for all enumeration types.

**Scope note.** `SolutionContinuity_t` captures *storage and DOF-sharing semantics* only—it answers “are interface DOFs shared or element-local?” It does not encode the polynomial basis, smoothness order, patch topology, or support-radius metadata that isogeometric analysis (IGA), meshless, or other non-standard discretizations may additionally require. Those richer descriptions belong in `SolutionInterpolation_t` (CPEX 0045) or in future extensions; this CPEX is intentionally scoped to the coarser, universally applicable continuous/discontinuous distinction.

### 3.3 SIDS Node Definition

The continuity is stored as a new optional child node of `FlowSolution_t`. The amended SIDS definition adds one line (shown with the + marker). The definition below includes the `SolutionInterpolation_t` child introduced by CPEX 0045:

#### `FlowSolution_t` (amended, SIDS Section 7.7)

```
FlowSolution_t< int CellDimension, int IndexDimension,
int VertexSize[IndexDimension],
int CellSize[IndexDimension] > :=
{
List( Descriptor_t Descriptor1 ... DescriptorN ) ;          (o)
GridLocation_t GridLocation ;                               (o/d)
Rind_t<IndexDimension> Rind ;                               (o/d)
IndexRange_t<IndexDimension> PointRange ;                 (o)
IndexArray_t<IndexDimension, ListLength[], int> PointList ; (o)
List( DataArray_t<DataType, IndexDimension, DataSize[]>
DataArray1 ... DataArrayN ) ;                             (o)
```

```

DataClass_t DataClass ; (o) 196
197
DimensionalUnits_t DimensionalUnits ; (o) 198
199
SolutionInterpolation_t SolutionInterpolation ; (o) (CPEX 200
0045) 201
202
+ SolutionContinuity_t SolutionContinuity ; (o) 203
204
List( UserDefinedData_t UserDefinedData1 ... UserDefinedDataN ) ; (o) 205
} 206

```

The `SolutionContinuity` child is **optional** (o). When absent, the continuity is unspecified (`SolutionContinuityNull`), preserving full backward compatibility. 207  
208

The same optional child node is added to `DiscreteData_t` (SIDS Section 12.4), using identical structure and semantics: 209  
210

#### DiscreteData\_t (amended, SIDS Section 12.4) 211

```

DiscreteData_t< int CellDimension, int IndexDimension, 213
int VertexSize[IndexDimension], 214
int CellSize[IndexDimension] > := 215
{ 216
... (existing children unchanged) ... 217
218
+ SolutionContinuity_t SolutionContinuity ; (o) 219
} 220

```

The optional child node is also added to `ZoneSubRegion_t` (SIDS Section 7.9). `ZoneSubRegion_t` holds `DataArray_t` nodes for localized solutions, boundary fluxes, or fields on sub-regions where the continuous/discontinuous distinction applies equally: 221  
222  
223

#### ZoneSubRegion\_t (amended, SIDS Section 7.9) 224

```

ZoneSubRegion_t< int RegionCellDimension > := 226
{ 227
... (existing children unchanged) ... 228
229
+ SolutionContinuity_t SolutionContinuity ; (o) 230
} 231

```

### 3.3.1 Node Placement Rationale 232

The node is placed as a child of `FlowSolution_t` (rather than, for example, `Zone_t` or `SolutionInterpolation_t`) because: 233  
234

- The continuity is a property of the *solution field*, not of the mesh or the zone. 235
- A single zone may contain multiple `FlowSolution_t` nodes with different continuity—for example, an HDG solver storing both volume (discontinuous) and trace (continuous) solutions, or a multi-physics code storing a continuous displacement field alongside a discontinuous stress field. 236  
237  
238  
239
- This parallels the existing `GridLocation_t` child of `FlowSolution_t`, which is also a per-solution-node property. 240  
241
- Placing it within the CPEX 0045 `SolutionInterpolation_t` would incorrectly limit the scope to higher-order solutions, when the distinction applies at any order. 242  
243

### 3.4 New MLL Functions 244

#### 3.4.1 C API 245

```

1  /*
2  * Write the solution continuity for a FlowSolution_t node.
3  *
4  * fn   - file number
5  * B    - base index
6  * Z    - zone index
7  * S    - solution index
8  * type - SolutionContinuity_t value
9  */
10 int cg_sol_continuity_write(
11     int fn, int B, int Z, int S,
12     CGNS_ENUMT(SolutionContinuity_t) type);
13
14 /*
15 * Read the solution continuity for a FlowSolution_t node.
16 *
17 * fn   - file number
18 * B    - base index
19 * Z    - zone index
20 * S    - solution index
21 * type - [out] SolutionContinuity_t value;
22 *       returns SolutionContinuityNull if node absent
23 */
24 int cg_sol_continuity_read(
25     int fn, int B, int Z, int S,
26     CGNS_ENUMT(SolutionContinuity_t) *type);
27
28 /*
29 * Write/read the solution continuity for a DiscreteData_t node.
30 * Same semantics as the FlowSolution_t counterparts.
31 *
32 * fn   - file number
33 * B    - base index
34 * Z    - zone index
35 * D    - discrete data index
36 * type - SolutionContinuity_t value
37 */
38 int cg_discrete_continuity_write(
39     int fn, int B, int Z, int D,
40     CGNS_ENUMT(SolutionContinuity_t) type);
41
42 int cg_discrete_continuity_read(
43     int fn, int B, int Z, int D,
44     CGNS_ENUMT(SolutionContinuity_t) *type);
45
46 /*
47 * Write/read the solution continuity for a ZoneSubRegion_t node.

```

```

48  * Same semantics as the FlowSolution_t counterparts.
49  *
50  * fn   - file number
51  * B    - base index
52  * Z    - zone index
53  * SR   - sub-region index
54  * type - SolutionContinuity_t value
55  */
56  int cg_subreg_continuity_write(
57      int fn, int B, int Z, int SR,
58      CGNS_ENUMT(SolutionContinuity_t) type);
59
60  int cg_subreg_continuity_read(
61      int fn, int B, int Z, int SR,
62      CGNS_ENUMT(SolutionContinuity_t) *type);
63
64  /*
65  * Return the string name for a SolutionContinuity_t value.
66  * Matches the convention of cg_GridLocationName, cg_DataTypeName, etc.
67  * Returns NULL for out-of-range values.
68  */
69  CGNSDLL const char *cg_SolutionContinuityName(
70      CGNS_ENUMT(SolutionContinuity_t) type);

```

**Behavior:**

246

- The `_write` functions create (or overwrite) the `SolutionContinuity` child node under the specified `FlowSolution_t`, `DiscreteData_t`, or `ZoneSubRegion_t` node. Writing `SolutionContinuityNull` removes the node if present. 247  
248  
249
- The `_read` functions return the value stored in the `SolutionContinuity` child node. If the node does not exist (legacy file or unset), they return `SolutionContinuityNull` without error. 250  
251

**3.4.2 Fortran API**

252

```

1  ! FlowSolution_t
2  subroutine cg_sol_continuity_write_f( &
3      fn, B, Z, S, continuity, ier)
4      integer, intent(in) :: fn, B, Z, S, continuity
5      integer, intent(out) :: ier
6
7  subroutine cg_sol_continuity_read_f( &
8      fn, B, Z, S, continuity, ier)
9      integer, intent(in) :: fn, B, Z, S
10     integer, intent(out) :: continuity, ier
11
12  ! DiscreteData_t
13  subroutine cg_discrete_continuity_write_f( &
14      fn, B, Z, D, continuity, ier)
15     integer, intent(in) :: fn, B, Z, D, continuity
16     integer, intent(out) :: ier
17
18  subroutine cg_discrete_continuity_read_f( &
19      fn, B, Z, D, continuity, ier)
20     integer, intent(in) :: fn, B, Z, D
21     integer, intent(out) :: continuity, ier
22
23  ! ZoneSubRegion_t
24  subroutine cg_subreg_continuity_write_f( &
25      fn, B, Z, SR, continuity, ier)
26     integer, intent(in) :: fn, B, Z, SR, continuity
27     integer, intent(out) :: ier
28
29  subroutine cg_subreg_continuity_read_f( &

```

```

30     fn, B, Z, SR, continuity, ier)
31     integer, intent(in) :: fn, B, Z, SR
32     integer, intent(out) :: continuity, ier

```

### 3.5 Internal Data Structure Changes

253

The `cgns_sol` structure (`FlowSolution_t`), `cgns_discrete` structure (`DiscreteData_t`), and `cgns_subreg` structure (`ZoneSubRegion_t`) in `cgns_header.h` each gain two fields:

254

255

```

1  typedef struct {          /* FlowSolution_t node */
2      /* ... existing fields ... */
3
4      /* CPEX 0050 */
5      CGNS_ENUMT(SolutionContinuity_t) continuity;
6      double continuity_id; /* CGIO node id for overwrite/delete */
7  } cgns_sol;
8
9  typedef struct {          /* DiscreteData_t node */
10     /* ... existing fields ... */
11
12     /* CPEX 0050 */
13     CGNS_ENUMT(SolutionContinuity_t) continuity;
14     double continuity_id;
15 } cgns_discrete;
16
17 typedef struct {          /* ZoneSubRegion_t node */
18     /* ... existing fields ... */
19
20     /* CPEX 0050 */
21     CGNS_ENUMT(SolutionContinuity_t) continuity;
22     double continuity_id;
23 } cgns_subreg;

```

The `continuity` field defaults to `SolutionContinuityNull` (value 0), which is the zero-initialization value and the “not specified” state. The `continuity_id` field stores the CGIO node identifier so that the write function can overwrite or delete the child node; it defaults to 0 (no node).

256

257

258

### 3.6 Internal Lookup Table Initialization

259

The string-to-enum mapping table must be initialized in `cgns_internals.c` so that `cgi_SolutionContinuity()` can resolve the C1 character data stored in CGNS files:

260

261

```

1  const char *SolutionContinuityName[NofValidSolutionContinuity] = {
2      "Null",
3      "UserDefined",
4      "Continuous",
5      "Discontinuous"
6  };

```

### 3.7 MLL Implementation Requirements

262

#### 3.7.1 Internal Tree Parser Update (`cgns_internals.c`)

263

When a CGNS file is opened, the MLL tree parser crawls the HDF5/CGIO tree and populates internal C structures. The parser must be updated to recognize child nodes with label `SolutionContinuity_t` during the traversal of `FlowSolution_t` (`cgi_read_sol`), `DiscreteData_t` (`cgi_read_discrete`), and `ZoneSubRegion_t` (`cgi_read_subregion`).

264

265

266

267

The required behavior is:

1. During memory allocation of `cgns_sol`, `cgns_discrete`, and `cgns_subreg`, the `continuity` field is zero-initialized, which corresponds to `SolutionContinuityNull`.
2. During the child-node traversal, if a child with label "SolutionContinuity\_t" is found, read its C1 string data and map it to the enum value via `cgi_SolutionContinuity(string_data, &continuity)`. Store the result in the struct's `continuity` field.
3. If no such child node exists, the field retains its default `SolutionContinuityNull` value.

This follows the existing pattern used for `GridLocation_t` parsing in `cgi_read_sol`.

### 3.7.2 Read Function Error Handling

The `_read` functions (e.g., `cg_sol_continuity_read`) return the cached struct field rather than performing a fresh HDF5 read. Because the internal tree parser already handles the absent-node case by leaving the field at `SolutionContinuityNull`, the read function simply returns the struct value and `CG_OK`:

```

1 int cg_sol_continuity_read(int fn, int B, int Z, int S,
2                           CGNS_ENUMT(SolutionContinuity_t) *type) {
3     cgns_sol *sol;
4
5     cg = cgi_get_file(fn);
6     if (cg == 0) return CG_ERROR;
7
8     if (cgi_check_mode(cg->filename, cg->mode, CG_MODE_READ))
9         return CG_ERROR;
10
11    sol = cgi_get_sol(cg, B, Z, S);
12    if (sol == NULL) return CG_ERROR;
13
14    *type = sol->continuity;
15
16    return CG_OK;
17 }

```

This architecture avoids the need to intercept `CG_NODE_NOT_FOUND` errors at read time—the absent-node case is handled once during file open, and all subsequent reads are served from the in-memory cache. Legacy code that iterates over solutions and queries the continuity will always receive `CG_OK` with `SolutionContinuityNull` for files that predate this CPEX.

### 3.7.3 Write Function Implementation

The `_write` functions create (or overwrite) the `SolutionContinuity` child node. Writing `SolutionContinuityNull` removes the node if present, restoring the “unspecified” state. The reference implementation for `FlowSolution_t` is shown below; the `DiscreteData_t` and `ZoneSubRegion_t` variants are analogous (substituting the appropriate internal structure accessor).

```

1  int cg_sol_continuity_write(int fn, int B, int Z, int S,
2                               CGNS_ENUMT(SolutionContinuity_t) type) {
3      cgns_sol *sol;
4      cgsizes_t length;
5      double posit_id;
6
7      cg = cgi_get_file(fn);
8      if (cg == 0) return CG_ERROR;
9
10     if (cgi_check_mode(cg->filename, cg->mode, CG_MODE_WRITE))
11         return CG_ERROR;
12
13     sol = cgi_get_sol(cg, B, Z, S);
14     if (sol == NULL) return CG_ERROR;
15
16     /* Validate enum value */
17     switch (type) {
18         case CGNS_ENUMV(SolutionContinuityNull):
19         case CGNS_ENUMV(SolutionContinuityUserDefined):
20         case CGNS_ENUMV(SolutionContinuous):
21         case CGNS_ENUMV(SolutionDiscontinuous):
22         break;
23     default:
24         cgi_error("Invalid SolutionContinuity value: %d", (int)type);
25         return CG_ERROR;
26     }
27
28     /* Delete existing node if writing Null */
29     if (type == CGNS_ENUMV(SolutionContinuityNull)) {
30         if (sol->continuity !=
31             CGNS_ENUMV(SolutionContinuityNull)) {
32             if (cgi_delete_node(sol->id,
33                                 sol->continuity_id))
34                 return CG_ERROR;
35             sol->continuity =
36                 CGNS_ENUMV(SolutionContinuityNull);
37             sol->continuity_id = 0;
38         }
39         return CG_OK;
40     }
41
42     /* Delete existing node before overwrite */
43     if (sol->continuity !=
44         CGNS_ENUMV(SolutionContinuityNull)) {
45         if (cgi_delete_node(sol->id,
46                             sol->continuity_id))
47             return CG_ERROR;
48         sol->continuity_id = 0;
49     }
50
51     /* Create the child node */
52     length = (cgsizes_t)strlen(
53         SolutionContinuityName[type]);
54     if (cgi_new_node(sol->id, "SolutionContinuity",
55                     "SolutionContinuity_t", &posit_id,
56                     "C1", 1, &length,
57                     (void *)SolutionContinuityName[type]))
58         return CG_ERROR;
59
60     sol->continuity = type;
61     sol->continuity_id = posit_id;
62     return CG_OK;
63 }

```

**Note:** The `continuity_id` field (a double storing the CGIO node identifier) must be added to the 290

internal structures alongside the `continuity` field so that the node can be located for overwrite or deletion. This follows the pattern used by other optional child nodes in the MLL (e.g., `GridLocation` tracking via `location_id` fields in some internal paths).

### 3.8 Data Layout for Discontinuous Solutions at Shared Mesh Vertices

A common and important case arises when a DG solver uses a mesh whose vertices are shared in the standard CGNS sense—enabling mesh walking, stream tracing, and connectivity queries via the standard `Elements_t` tables—but stores element-local, independent solution values at each element’s copy of those vertices. At a shared vertex adjacent to  $k$  elements there are  $k$  distinct solution values, one per element; readers must not average or conflate them.

The correct CGNS encoding for this pattern is `GridLocation_t = Vertex` combined with `SolutionContinuity_t = SolutionDiscontinuous` and a `PointList` that lists every element-vertex pair explicitly (repeated vertex indices are expected and meaningful). The `SolutionContinuity_t` node is the machine-readable flag that tells readers these repeated entries represent independent DOFs, not duplicate data.

Without the `SolutionContinuity_t` node, a reader cannot distinguish this layout from a malformed continuous solution with duplicate entries, which is why the node is necessary. With it, a post-processor or coupling framework can correctly:

- walk the mesh using the shared-vertex connectivity for geometry operations (stream tracing, surface intersection, etc.);
- read solution values element-locally, preserving the jumps at shared interfaces.

For element-local DOFs at non-vertex positions (interior quadrature points, face quadrature points, etc.), use `GridLocation_t = InterpolationPoints` with `SolutionInterpolation_t` metadata from CPEX 0045. `GridLocation_t = Vertex` with a `PointList` is specifically the correct choice when the DOF positions coincide with mesh vertices and the shared-vertex mesh connectivity is to be preserved.

`GridLocation_t = Vertex` with `SolutionDiscontinuous` and *no* `PointList` is not valid: the SIDS `DataSize` formula for `Vertex` yields `DataSize = VertexSize` (the shared-vertex count), which cannot represent the larger element-local DOF count. Similarly, `PointRange` is not suitable for this case because a contiguous range of vertex indices cannot express the element-local repetition of shared vertices; only `PointList` provides the necessary per-element DOF ordering. `GridLocation_t = CellCenter` with `SolutionDiscontinuous` remains valid since cell-centered data is inherently one value per element with no shared-entity ambiguity.

This approach extends naturally to mixed-element meshes (`MIXED`, `NGON_n`, `NFACE_n`). The `PointList` simply lists vertex indices for each element in element order; elements with different vertex counts contribute different numbers of entries. The `SolutionContinuity_t = SolutionDiscontinuous` flag tells readers that repeated vertex indices represent independent DOFs regardless of element type heterogeneity.

**SIDS Amendment (PointList semantics under Discontinuous solutions):** When `SolutionContinuity_t = SolutionDiscontinuous` and a `PointList` is present under `FlowSolution_t` with `GridLocation_t = Vertex`, repeated vertex indices are permitted and represent element-local DOFs. The DOF-to-element mapping is defined by traversal of the zone’s `Elements_t` sections in ascending section index order, then by each section’s connectivity

array order. Solution array position  $i$  corresponds to the  $i$ -th entry of the `PointList` under this traversal. This is an explicit amendment to the `PointList` description in SIDS §7.7 for this case; repeated indices under a continuous solution remain undefined and invalid.

**Note:** For element-local DOFs at non-vertex positions, or when the implicit element-ordering covenant above is inconvenient, `GridLocation_t = InterpolationPoints` with `SolutionInterpolation_t` from CPEX 0045 is the preferred encoding and does not require repeated indices.

See Example 6.7 for the complete C encoding of this pattern.

### 3.9 Interaction with `Rind_t` (Ghost Layers)

`SolutionContinuity_t` applies uniformly to the entire solution field, including any rind (ghost) entries declared by `Rind_t`. A `SolutionDiscontinuous` label means that all data entries—interior and rind alike—are element-local: each rind entry belongs to one element and must not be interpreted as a shared vertex value. DG solvers that use ghost layers for MPI halo exchange may therefore store element-local rind DOFs in the standard way; the `SolutionDiscontinuous` label covers them. Rind entries of a `SolutionContinuous` solution retain their usual shared-DOF semantics. No special treatment of the rind is required beyond what `Rind_t` already specifies; the continuity label is consistent across the full data array.

## 4 Interaction with CPEX 0045 (Higher-Order Elements)

CPEX 0045 introduced `GridLocation_t = InterpolationPoints`, `SolutionInterpolation_t`, and the concept of per-element-type polynomial basis definitions. The present CPEX complements rather than modifies that infrastructure:

- CPEX 0045 answers: “What polynomial basis and order describes the solution within each element?”
- This CPEX answers: “Are the resulting DOFs shared across element interfaces or element-local?”

Both pieces of metadata are needed for a complete description of a higher-order solution field. They are orthogonal:

	Continuous	Discontinuous
Lagrange $p = 1$	CG FEM	DG $p = 1$
Lagrange $p = 3$	CG spectral	DG spectral
Hierarchical $p = 2$	CG (modal)	DG (modal)
$p = 0$ (constant)	—	FV / DG $p = 0$

For higher-order solutions using `GridLocation_t = InterpolationPoints`, the recommended practice is:

1. Use CPEX 0045 to define the element type, interpolation basis, and polynomial order via `SolutionInterpolation_t`.
2. Use this CPEX to set `SolutionContinuity_t` on the `FlowSolution_t` node to indicate whether the field is continuous or discontinuous.

For lower-order solutions using `GridLocation_t = Vertex` or `CellCenter`, this CPEX can be used independently of CPEX 0045. See Section 3.8 for the prescribed data layout when DG DOFs coincide with shared mesh vertices.

## 5 Backward Compatibility

- **Existing files:** Fully compatible. The `SolutionContinuity` node is optional; its absence is equivalent to `SolutionContinuityNull`. No existing file is invalidated.
- **Existing readers:** Unaffected. Per the CGNS convention, readers silently ignore unknown child nodes. An older library reading a file with `SolutionContinuity` nodes will simply skip them.
- **Existing writers:** Unaffected. Codes that do not call the new API functions produce files without `SolutionContinuity` nodes, identical to the current behavior.
- **ABI:** The `cgns_sol`, `cgns_discrete`, and `cgns_subreg` structures each grow by two fields (one `int`-sized enum and one `double` node identifier). This is an ABI change requiring recompilation of code that directly accesses the internal structures. The public API is purely additive (new functions, new enum), so no existing function signatures change.
- **Parallel I/O (`cgp_*`):** The `SolutionContinuity` metadata node is a small scalar child of `FlowSolution_t` and carries no per-element data. It is written and read via the standard collective-write / independent-read model, identical to how `GridLocation` is handled in the parallel MLL. No new `cgp_*` variants are required.

## 6 Examples

### 6.1 Example 1: Writing a Continuous Galerkin Solution (C)

```

1  /* Create a vertex-located solution node */
2  int S;
3  cg_sol_write(fn, B, Z, "FlowSolution", CGNS_ENUMV(Vertex), &S);
4
5  /* Mark as continuous (CG FEM) */
6  cg_sol_continuity_write(fn, B, Z, S,
7      CGNS_ENUMV(SolutionContinuous));
8
9  /* Write solution data (one value per mesh vertex, shared) */
10 cg_field_write(fn, B, Z, S, CGNS_ENUMV(RealDouble),
11     "Pressure", pressure, &F);

```

### 6.2 Example 2: Writing a Discontinuous Galerkin Solution (C)

```

1  /* Create an interpolation-point-located solution node */
2  int S;
3  cg_sol_write(fn, B, Z, "FlowSolution_DG",
4      CGNS_ENUMV(InterpolationPoints), &S);
5
6  /* Set polynomial order (CPEX 0045) */
7  cg_sol_interpolation_order_write(fn, B, Z, S, 3, 0);
8
9  /* Mark as discontinuous */
10 cg_sol_continuity_write(fn, B, Z, S,

```

```

11     CGNS_ENUMV(SolutionDiscontinuous));
12
13  /* Write element-local DOFs */
14  cg_field_write(fn, B, Z, S, CGNS_ENUMV(RealDouble),
15               "Pressure", pressure_dg, &F);

```

### 6.3 Example 3: Reading and Dispatching on Continuity (C)

388

```

1  CGNS_ENUMT(SolutionContinuity_t) repr;
2  cg_sol_continuity_read(fn, B, Z, S, &repr);
3
4  switch (repr) {
5  case CGNS_ENUMV(SolutionContinuous):
6      /* Shared DOFs: interpolate using global mesh connectivity */
7      visualize_continuous(fn, B, Z, S);
8      break;
9  case CGNS_ENUMV(SolutionDiscontinuous):
10     /* Element-local DOFs: render per-element, preserve jumps */
11     visualize_discontinuous(fn, B, Z, S);
12     break;
13 default:
14     /* Null or UserDefined: fall back to legacy heuristics */
15     visualize_legacy(fn, B, Z, S);
16     break;
17 }

```

### 6.4 Example 4: HDG Solver with Both Continuity Types (C)

389

```

1  /* Volume DOFs (discontinuous) */
2  int S_vol;
3  cg_sol_write(fn, B, Z, "VolumeSolution",
4             CGNS_ENUMV(InterpolationPoints), &S_vol);
5  cg_sol_continuity_write(fn, B, Z, S_vol,
6                        CGNS_ENUMV(SolutionDiscontinuous));
7
8  /* Trace DOFs on element faces (continuous on skeleton).
9   * Case A: p=0 traces (one DOF per face) -- use FaceCenter. */
10 int S_trace;
11 cg_sol_write(fn, B, Z, "TraceSolution",
12            CGNS_ENUMV(FaceCenter), &S_trace);
13 cg_sol_continuity_write(fn, B, Z, S_trace,
14                       CGNS_ENUMV(SolutionContinuous));
15
16 /* Case B: higher-order traces -- use InterpolationPoints
17  * with CPEX 0045 metadata for the face polynomial basis. */
18 int S_trace_ho;
19 cg_sol_write(fn, B, Z, "TraceSolution_HO",
20            CGNS_ENUMV(InterpolationPoints), &S_trace_ho);
21 cg_sol_interpolation_order_write(fn, B, Z, S_trace_ho, 2, 0);
22 cg_sol_continuity_write(fn, B, Z, S_trace_ho,
23                       CGNS_ENUMV(SolutionContinuous));

```

### 6.5 Example 5: Fortran Usage

390

```

1 integer :: S, ier, repr
2
3 ! Write a continuous solution
4 call cg_sol_write_f(fn, B, Z, 'FlowSolution', Vertex, S, ier)
5 call cg_sol_continuity_write_f(fn, B, Z, S, &
6     SolutionContinuous, ier)
7
8 ! Read continuity
9 call cg_sol_continuity_read_f(fn, B, Z, S, repr, ier)
10 if (repr == SolutionDiscontinuous) then
11     ! Element-local interpolation
12 end if

```

## 6.6 Example 6: Low-Order Finite Volume (DG $p = 0$ )

391

```

1 /* Cell-centered FV solution */
2 int S;
3 cg_sol_write(fn, B, Z, "FlowSolution",
4     CGNS_ENUMV(CellCenter), &S);
5
6 /* Explicitly mark as discontinuous */
7 cg_sol_continuity_write(fn, B, Z, S,
8     CGNS_ENUMV(SolutionDiscontinuous));

```

This example illustrates that continuity metadata provides value even for traditional low-order methods. Although a `CellCenter` solution is inherently element-local, explicitly recording the continuity assists post-processors and coupling frameworks in programmatically distinguishing between solution types.

392

393

394

395

## 6.7 Example 7: DG Solution on a Shared-Vertex Mesh (C)

396

This example addresses the case where the mesh uses shared vertices (enabling stream tracing and mesh walking via standard `Elements_t` connectivity), but each element stores its own independent solution values at those vertices. The `SolutionContinuity_t` node is the flag that tells readers the repeated vertex indices in the `PointList` are independent DOFs, not duplicates.

397

398

399

400

```

1 /*
2  * Mesh: nVerts shared vertices, nElems elements,
3  *       nVertsPerElem vertices per element.
4  * Solution: element-local DOFs at each element's vertices.
5  *       Total DOFs = nElems * nVertsPerElem.
6  */
7
8 /* Build the point list: for each element, list its vertex indices.
9  * Shared vertices appear multiple times -- once per adjacent element. */
10 cgsized_t point_list[nElems * nVertsPerElem];
11 for (int e = 0; e < nElems; e++)
12     for (int v = 0; v < nVertsPerElem; v++)
13         point_list[e * nVertsPerElem + v] =
14             element_connectivity[e][v]; /* 1-based vertex index */
15
16 /* Create solution node with PointList using the idiomatic
17  * cg_sol_ptset_write: Vertex location preserves shared-mesh
18  * connectivity for geometry operations.
19  * DataSize = nElems * nVertsPerElem (not nVerts). */
20 int S;
21 cg_sol_ptset_write(fn, B, Z, "FlowSolution_DG",
22     CGNS_ENUMV(Vertex),
23     CGNS_ENUMV(PointList),

```

```

24     (cgsize_t)(nElems * nVertsPerElem),
25     point_list, &S);
26
27  /* Mark as discontinuous: values at repeated vertex indices
28   * are independent across elements, ordered by Elements_t
29   * section index then connectivity order. */
30  cg_sol_continuity_write(fn, B, Z, S,
31     CGNS_ENUMV(SolutionDiscontinuous));
32
33  /* Write element-local solution data */
34  cg_field_write(fn, B, Z, S, CGNS_ENUMV(RealDouble),
35     "Pressure", pressure_dg, &F);

```

A reader distinguishes this from a continuous solution as follows:

401

```

1  CGNS_ENUMT(SolutionContinuity_t) cont;
2  cg_sol_continuity_read(fn, B, Z, S, &cont);
3
4  if (cont == CGNS_ENUMV(SolutionDiscontinuous)) {
5     /* PointList entries are element-local DOF indices.
6      * Values at the same vertex index from different elements
7      * are independent -- do not average. */
8  }

```

## 6.8 Example 8: DiscreteData\_t with Discontinuous Error Indicator (C)

402

This example shows a DG error indicator stored in a DiscreteData\_t node. The usage mirrors FlowSolution\_t; the only difference is the API prefix (cg\_discrete\_ instead of cg\_sol\_).

403

404

```

1  /* Create a cell-centered discrete data node for error indicators */
2  int D;
3  cg_discrete_write(fn, B, Z, "ErrorIndicator", &D);
4  cg_goto(fn, B, "Zone_t", Z, "DiscreteData_t", D, "end");
5  cg_gridlocation_write(CGNS_ENUMV(CellCenter));
6
7  /* Mark as discontinuous (element-local error estimates) */
8  cg_discrete_continuity_write(fn, B, Z, D,
9     CGNS_ENUMV(SolutionDiscontinuous));
10
11 /* Write per-element error data */
12 cg_goto(fn, B, "Zone_t", Z, "DiscreteData_t", D, "end");
13 cg_array_write("ElementError", CGNS_ENUMV(RealDouble),
14     1, &nElems, error_data);

```

## 6.9 Example 9: ZoneSubRegion\_t with Boundary Flux Data (C)

405

This example stores discontinuous boundary flux data on a ZoneSubRegion\_t that references a boundary condition region.

406

407

```

1  /* Create a sub-region referencing a BC */
2  int SR;
3  /* CellDimension is the zone's cell dimension (e.g., 3 for a 3-D volume
4   * zone), so CellDimension - 1 is the face dimension. */
5  cg_subreg_bcname_write(fn, B, Z, "WallFluxes",
6     CellDimension - 1, "Wall_BC", &SR);
7
8  /* Mark as discontinuous (each face owns its own flux values) */
9  cg_subreg_continuity_write(fn, B, Z, SR,
10     CGNS_ENUMV(SolutionDiscontinuous));
11

```

```

12 /* Write flux data on the sub-region */
13 cg_goto(fn, B, "Zone_t", Z, "ZoneSubRegion_t", SR, "end");
14 cg_array_write("HeatFlux", CGNS_ENUMV(RealDouble),
15               1, &nBCFaces, heat_flux);

```

## 6.10 Example 10: Solver Restart with Continuity-Aware Read (C)

408

A solver reading a restart file must reconstruct the solution state correctly based on the continuity. This example shows how a restart reader uses the continuity metadata to dispatch between shared-DOF and element-local reconstruction.

409

410

411

```

1 /* Read solution metadata */
2 char sol_name[CGIO_MAX_NAME_LENGTH+1];
3 CGNS_ENUMT(GridLocation_t) loc;
4 cg_sol_info(fn, B, Z, S, sol_name, &loc);
5
6 CGNS_ENUMT(SolutionContinuity_t) cont;
7 cg_sol_continuity_read(fn, B, Z, S, &cont);
8
9 /* Read the field data. In a generic restart loop, call
10 * cg_field_info(fn,B,Z,S,F,...) first to discover field names
11 * and types before calling cg_field_read. */
12 cg_field_read(fn, B, Z, S, "Pressure",
13              CGNS_ENUMV(RealDouble), range_min, range_max,
14              pressure_buf);
15
16 /* Reconstruct based on continuity */
17 if (cont == CGNS_ENUMV(SolutionContinuous)) {
18     /* DOFs are shared at mesh vertices/edges/faces.
19     * Map directly to the global DOF vector. */
20     reconstruct_continuous(pressure_buf, loc, mesh);
21 } else if (cont == CGNS_ENUMV(SolutionDiscontinuous)) {
22     /* DOFs are element-local. Each element owns its
23     * own copy of values at shared interfaces. */
24     reconstruct_discontinuous(pressure_buf, loc, mesh);
25 } else {
26     /* Null: legacy file. Infer from data size or
27     * assume continuous for backward compatibility. */
28     reconstruct_legacy(pressure_buf, loc, mesh);
29 }

```

## 6.11 Example 11: FSI Coupling — Structural and Fluid Solutions (C)

412

In Fluid-Structure Interaction (FSI) applications, a single zone may carry both fluid and structural solution fields with different continuity types. A continuous FEM structural solver stores shared displacement DOFs, while a DG fluid solver stores element-local flow variables. Separate FlowSolution\_t nodes with distinct continuity metadata allow coupling frameworks to correctly interpret and transfer each field.

413

414

415

416

417

```

1 /* Structural solution: continuous FEM displacements */
2 int S_struct;
3 cg_sol_write(fn, B, Z, "StructuralDisplacement",
4             CGNS_ENUMV(Vertex), &S_struct);
5 cg_sol_continuity_write(fn, B, Z, S_struct,
6                        CGNS_ENUMV(SolutionContinuous));
7 cg_field_write(fn, B, Z, S_struct, CGNS_ENUMV(RealDouble),
8               "DisplacementX", disp_x, &F);
9 cg_field_write(fn, B, Z, S_struct, CGNS_ENUMV(RealDouble),
10              "DisplacementY", disp_y, &F);
11

```

```

12 /* Fluid solution: discontinuous DG flow field */
13 int S_fluid;
14 cg_sol_write(fn, B, Z, "FluidSolution",
15             CGNS_ENUMV(InterpolationPoints), &S_fluid);
16 cg_sol_interpolation_order_write(fn, B, Z, S_fluid, 3, 0);
17 cg_sol_continuity_write(fn, B, Z, S_fluid,
18                       CGNS_ENUMV(SolutionDiscontinuous));
19 cg_field_write(fn, B, Z, S_fluid, CGNS_ENUMV(RealDouble),
20              "Pressure", pressure, &F);
21 cg_field_write(fn, B, Z, S_fluid, CGNS_ENUMV(RealDouble),
22              "VelocityX", vel_x, &F);
23
24 /* A coupling framework reading this zone can query
25  * each FlowSolution_t's continuity to determine the
26  * correct interpolation and transfer strategy:
27  *   - Continuous fields: global shared-DOF interpolation
28  *   - Discontinuous fields: element-local projection */

```

The same pattern applies when the structural and fluid domains are in separate zones. The continuity metadata on each zone's `FlowSolution_t` nodes enables the coupling framework to apply the correct inter-code transfer operators without relying on naming conventions or external configuration.

## 6.12 Example 12: Conjugate Heat Transfer — Perfect and Imperfect Thermal Contact (C)

Conjugate heat transfer (CHT) simulations couple a fluid temperature field (often stored at cell centers or DG DOFs) with a solid temperature field (often stored as a continuous CG FEM solution). The continuity metadata disambiguates the two layouts and informs the coupling framework how to perform the interface transfer.

```

1 /* ---- Fluid zone: FV/DG temperature (discontinuous) ---- */
2 int S_fluid;
3 cg_sol_write(fn, B, fluidZone, "ThermalSolution",
4             CGNS_ENUMV(CellCenter), &S_fluid);
5 cg_sol_continuity_write(fn, B, fluidZone, S_fluid,
6                       CGNS_ENUMV(SolutionDiscontinuous));
7 cg_field_write(fn, B, fluidZone, S_fluid,
8               CGNS_ENUMV(RealDouble), "Temperature", T_fluid, &F);
9
10 /* ---- Solid zone: CG FEM temperature (continuous) ---- */
11 int S_solid;
12 cg_sol_write(fn, B, solidZone, "ThermalSolution",
13             CGNS_ENUMV(Vertex), &S_solid);
14 cg_sol_continuity_write(fn, B, solidZone, S_solid,
15                       CGNS_ENUMV(SolutionContinuous));
16 cg_field_write(fn, B, solidZone, S_solid,
17               CGNS_ENUMV(RealDouble), "Temperature", T_solid, &F);
18
19 /* ---- Interface heat flux on the solid side (face-local) ---- */
20 int SR;
21 cg_subreg_bcname_write(fn, B, solidZone, "FluidInterface",
22                      CellDimension - 1, "FluidInterface_BC", &SR);
23 cg_subreg_continuity_write(fn, B, solidZone, SR,
24                           CGNS_ENUMV(SolutionDiscontinuous));
25 cg_goto(fn, B, "Zone_t", solidZone,
26         "ZoneSubRegion_t", SR, "end");
27 cg_array_write("HeatFlux", CGNS_ENUMV(RealDouble),
28               1, &nInterfaceFaces, heat_flux);

```

With *perfect* thermal contact the fluid-side and solid-side interface temperatures are equal; the coupling framework can enforce this constraint knowing the solid temperature is `SolutionContinuous`

(shared vertex DOFs) and the fluid temperature is `SolutionDiscontinuous` (cell-centered). 429

With *imperfect* thermal contact (finite thermal resistance), a temperature jump  $\Delta T$  exists at the 430  
interface. Each domain retains its own interface temperature, and the coupling framework applies 431  
the resistance relation  $q = \Delta T/R$  between them. No change to the `SolutionContinuity_t` labels 432  
is needed: the fluid and solid temperatures are stored and tagged independently in their respective 433  
zones, and the jump is an emergent result of the coupling rather than a property of either field's 434  
storage layout. 435

## 7 SIDS Impact 436

1. **New enumeration type:** `SolutionContinuity_t` is added to the SIDS building-block 437  
structure definitions (Section 4 of the SIDS document), alongside existing enumeration types 438  
such as `GridLocation_t` (Section 4.5) and `DataClass_t` (Section 4.1). 439
2. **FlowSolution\_t amendment:** The `FlowSolution_t` node definition (SIDS Section 7.7) is 440  
amended to include an optional `SolutionContinuity_t` child node. 441
3. **DiscreteData\_t amendment:** The `DiscreteData_t` node definition (SIDS Section 12.4) is 442  
similarly amended. 443
4. **ZoneSubRegion\_t amendment:** The `ZoneSubRegion_t` node definition (SIDS Section 7.9) 444  
is similarly amended. 445
5. **PointList semantics amendment:** The description of `PointList` under `FlowSolution_t` 446  
in SIDS §7.7 is amended to permit repeated vertex indices when `SolutionContinuity_t =` 447  
`SolutionDiscontinuous`. Repeated indices in that context represent element-local DOFs and 448  
are not malformed duplicates; the DOF-to-element mapping follows ascending `Elements_t` 449  
section index order, then connectivity order within each section. This amendment does not 450  
affect `PointList` semantics under any other condition. 451
6. No existing SIDS node types, enumerations, or conventions are modified or removed. 452

## 8 File Mapping Impact 453

The `SolutionContinuity` child node is stored identically to `GridLocation`: as a character node 454  
containing the string name of the enumeration value. The file-mapping attributes are: 455

Attribute	Value
Node name	<code>SolutionContinuity</code>
Label	<code>SolutionContinuity_t</code>
Data type	C1 (character) 456
Dimension	1
Dimension value	length of enumeration string
Data	enumeration name string

This follows the same pattern as the `GridLocation` node under `FlowSolution_t`: node name 457  
"GridLocation", label "GridLocation\_t", data type "C1", with the enumeration string as data. 458

The node is written only when the continuity is not `SolutionContinuityNull`. When absent, readers default to `SolutionContinuityNull`, paralleling how an absent `GridLocation` node defaults to `Vertex`.

## 9 Design Decisions

Several design questions were evaluated during the development of this proposal. Each is addressed below with rationale grounded in existing CGNS conventions and SIDS precedent.

1. **Scope: `BCData_t` is excluded.** `SolutionContinuity_t` is *not* added to `BCData_t`. Boundary condition data arrays in `BCData_t` are sized by the associated `PointList/PointRange` and `GridLocation_t`, which already determine the data layout unambiguously. In DG methods, boundary conditions are typically applied weakly through numerical fluxes; the element-local nature of the boundary data follows from the solution continuity on the parent `FlowSolution_t` node. Adding `SolutionContinuity_t` to `BCData_t` would be redundant and would introduce a potential inconsistency if the BC continuity differs from the solution it constrains. Should a future use case require per-BC continuity metadata (e.g., coupling frameworks imposing boundary data in a different continuity than the interior, or interface conditions in CHT simulations with imperfect thermal contact where a flux-jump boundary condition is prescribed), it can be proposed as a separate, narrowly scoped CPEX without affecting this proposal. The `ZoneSubRegion_t` extension introduced here partially addresses such cases: interface fluxes or surface-local fields that require a continuity label can be stored in a `ZoneSubRegion_t` referencing the relevant boundary condition (see Example 6.12).

2. **Per-solution-node granularity (not per-field).** The continuity applies at the `FlowSolution_t` level, meaning all `DataArray_t` children share the same continuity. This follows the precedent set by `GridLocation_t`, which is likewise a per-solution-node property: all fields within a `FlowSolution_t` node share the same grid location. Per-`DataArray_t` granularity was rejected because it would add significant API and metadata complexity with no demonstrated need. Solvers that genuinely mix continuity types within a single zone (e.g., HDG methods with continuous traces and discontinuous volumes) should use separate `FlowSolution_t` nodes, exactly as they would use separate nodes for fields at different `GridLocation_t` values.

3. **CellCenter + SolutionContinuous: permitted but semantically inconsistent.** The MLL does *not* enforce consistency between `SolutionContinuity_t` and `GridLocation_t`. A `CellCenter` solution is inherently element-local, making `SolutionContinuous` semantically contradictory, but the library shall not reject such combinations. This follows the established CGNS convention of being permissive in what is accepted: the standard documents correct usage; validation is delegated to application-level tools such as `cgnscheck`.

**Convention:** `SolutionContinuity_t = SolutionContinuous` with `GridLocation_t = CellCenter` should not be used. `cgnscheck` should issue a warning when this combination is encountered.

4. **Recommended `cgnscheck` validation rules.** The following consistency checks are recommended for `cgnscheck` and similar validation tools:

**Warning:** `SolutionContinuity_t = SolutionContinuous` with `GridLocation_t = CellCenter` — semantically inconsistent since cell-centered data is inherently element-

local (Design Decision 3 above). 501

**Error:** `SolutionContinuity_t = SolutionDiscontinuous` with `GridLocation_t = Vertex` and `no PointList` — the SIDS `DataSize` formula yields `VertexSize` (the shared-vertex count), which cannot represent the larger element-local DOF count (Section 3.8). 502  
503  
504  
505

**Warning:** `SolutionContinuity_t` value outside the range `[0, NofValidSolutionContinuity)` — invalid enumeration value. 506  
507

## 10 Summary of API Additions 508

Symbol	Description
<code>SolutionContinuity_t</code>	New enumeration type: <code>Null</code> , <code>UserDefined</code> , <code>SolutionContinuous</code> , <code>SolutionDiscontinuous</code> .
<code>SolutionContinuityName</code>	String table for the enumeration values.
<code>NofValidSolutionContinuity</code>	Count of valid enumeration values (4).
<code>cg_SolutionContinuityName</code>	C: return the string name for a <code>SolutionContinuity_t</code> value (mirrors <code>cg_GridLocationName</code> ).
<code>cg_sol_continuity_write</code>	C: write solution continuity for a <code>FlowSolution_t</code> node.
<code>cg_sol_continuity_read</code>	C: read solution continuity for a <code>FlowSolution_t</code> node.
<code>cg_discrete_continuity_write</code>	C: write solution continuity for a <code>DiscreteData_t</code> node.
<code>cg_discrete_continuity_read</code>	C: read solution continuity for a <code>DiscreteData_t</code> node.
<code>cg_subreg_continuity_write</code>	C: write solution continuity for a <code>ZoneSubRegion_t</code> node.
<code>cg_subreg_continuity_read</code>	C: read solution continuity for a <code>ZoneSubRegion_t</code> node.
<code>cg_sol_continuity_write_f</code>	Fortran: write solution continuity ( <code>FlowSolution_t</code> ).
<code>cg_sol_continuity_read_f</code>	Fortran: read solution continuity ( <code>FlowSolution_t</code> ).
<code>cg_discrete_continuity_write_f</code>	Fortran: write solution continuity ( <code>DiscreteData_t</code> ).
<code>cg_discrete_continuity_read_f</code>	Fortran: read solution continuity ( <code>DiscreteData_t</code> ).
<code>cg_subreg_continuity_write_f</code>	Fortran: write solution continuity ( <code>ZoneSubRegion_t</code> ).
<code>cg_subreg_continuity_read_f</code>	Fortran: read solution continuity ( <code>ZoneSubRegion_t</code> ).

---

## 11 References

---

- |   |            |
|---|------------|
| 1. CGNS Standard Interface Data Structures (SIDS),<br><a href="https://cgns.org/standard/SIDS/CGNS_SIDS.html">https://cgns.org/standard/SIDS/CGNS_SIDS.html</a> | 511<br>512 |
| 2. CGNS Mid-Level Library—Flow Solution,<br><a href="https://cgns.org/standard/MLL/api/c_api.html">https://cgns.org/standard/MLL/api/c_api.html</a>             | 513<br>514 |
| 3. CPEX 0045, “Polynomial Data and Curved Grid Elements,”<br><a href="https://github.com/CGNS/CGNS/issues/577">https://github.com/CGNS/CGNS/issues/577</a>      | 515<br>516 |
| 4. Hesthaven, J.S. and Warburton, T., <i>Nodal Discontinuous Galerkin Methods</i> , Springer, 2008.   | 517        |
| 5. Cockburn, B., Karniadakis, G.E., and Shu, C.-W., eds., <i>Discontinuous Galerkin Methods: Theory, Computation, and Applications</i> , Springer, 2000.        | 518<br>519 |