



Process Management Interface for Exascale (PMIx) Standard

Version 4.0 (Draft)

1H2019

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 4.0 (Draft).

Comments: Please provide comments on the PMIx Standard by filing issues on the document repository <https://github.com/pmix/pmix-standard/issues> or by sending them to the PMIx Community mailing list at <https://groups.google.com/forum/#!forum/pmix>. Comments should include the version of the PMIx standard you are commenting about, and the page, section, and line numbers that you are referencing. Please note that messages sent to the mailing list from an unsubscribed e-mail address will be ignored.

Copyright © 2018-2019 PMIx Standard Review Board.

Permission to copy without fee all or part of this material is granted, provided the PMIx Standard Review Board copyright notice and the title of this document appear, and notice is given that copying is by permission of PMIx Standard Review Board.

This page intentionally left blank

Contents

1. Introduction	1
1.1. Charter	2
1.2. PMIx Standard Overview	2
1.2.1. Who should use the standard?	2
1.2.2. What is defined in the standard?	3
1.2.3. What is <i>not</i> defined in the standard?	3
1.2.4. General Guidance for PMIx Users and Implementors	4
1.3. PMIx Architecture Overview	4
1.3.1. The PMIx Reference Implementation (PRI)	6
1.3.2. The PMIx Reference RunTime Environment (PRRTE)	7
1.4. Organization of this document	7
1.5. Version 1.0: June 12, 2015	8
1.6. Version 2.0: Sept. 2018	8
1.7. Version 2.1: Dec. 2018	9
1.8. Version 2.2: Jan 2019	10
1.9. Version 3.0: Dec. 2018	10
1.10. Version 3.1: Jan. 2019	11
1.11. Version 4.0: June 2019	12
2. PMIx Terms and Conventions	13
2.1. Notational Conventions	15
2.2. Semantics	16
2.3. Naming Conventions	17
2.4. Procedure Conventions	17
2.5. Standard vs Reference Implementation	18
3. Data Structures and Types	19
3.1. Constants	20
3.1.1. PMIx Error Constants	21

3.1.2.	Macros for use with PMIx constants	25
3.2.	Data Types	25
3.2.1.	Key Structure	26
3.2.2.	Namespace Structure	27
3.2.3.	Rank Structure	28
3.2.4.	Process Structure	28
3.2.5.	Process structure support macros	28
3.2.6.	Process State Structure	30
3.2.7.	Process Information Structure	31
3.2.8.	Process Information Structure support macros	32
3.2.9.	Scope of Put Data	33
3.2.10.	Range of Published Data	34
3.2.11.	Data Persistence Structure	34
3.2.12.	Data Array Structure	35
3.2.13.	Data array structure support macros	35
3.2.14.	Value Structure	36
3.2.15.	Value structure support macros	37
3.2.16.	Info Structure	41
3.2.17.	Info structure support macros	41
3.2.18.	Info Type Directives	44
3.2.19.	Info Directive support macros	45
3.2.20.	Job Allocation Directives	47
3.2.21.	IO Forwarding Channels	47
3.2.22.	Environmental Variable Structure	47
3.2.23.	Environmental variable support macros	48
3.2.24.	Lookup Returned Data Structure	49
3.2.25.	Lookup data structure support macros	49
3.2.26.	Application Structure	52
3.2.27.	App structure support macros	53
3.2.28.	Query Structure	54
3.2.29.	Query structure support macros	54
3.2.30.	Attribute registration structure	56
3.2.31.	Attribute registration structure support macros	57

3.2.32.	PMIx Group Directives	59
3.2.33.	Byte Object Type	59
3.2.34.	Byte object support macros	59
3.2.35.	Data Array Structure	61
3.2.36.	Data array support macros	61
3.2.37.	Argument Array Macros	62
3.2.38.	Set Environment Variable	66
3.3.	Generalized Data Types Used for Packing/Unpacking	67
3.4.	Reserved attributes	68
3.4.1.	Initialization attributes	69
3.4.2.	Tool-related attributes	69
3.4.3.	Identification attributes	70
3.4.4.	Programming model attributes	71
3.4.5.	UNIX socket rendezvous socket attributes	71
3.4.6.	TCP connection attributes	72
3.4.7.	Global Data Storage (GDS) attributes	72
3.4.8.	General process-level attributes	72
3.4.9.	Scratch directory attributes	73
3.4.10.	Relative Rank Descriptive Attributes	73
3.4.11.	Information retrieval attributes	74
3.4.12.	Information storage attributes	75
3.4.13.	Size information attributes	76
3.4.14.	Memory information attributes	77
3.4.15.	Topology information attributes	77
3.4.16.	Request-related attributes	78
3.4.17.	Server-to-PMIx library attributes	79
3.4.18.	Server-to-Client attributes	80
3.4.19.	Event handler registration and notification attributes	80
3.4.20.	Fault tolerance attributes	81
3.4.21.	Spawn attributes	82
3.4.22.	Query attributes	84
3.4.23.	Log attributes	85
3.4.24.	Debugger attributes	86

3.4.25.	Resource manager attributes	87
3.4.26.	Environment variable attributes	87
3.4.27.	Job Allocation attributes	88
3.4.28.	Job control attributes	89
3.4.29.	Monitoring attributes	90
3.4.30.	Security attributes	91
3.4.31.	IO Forwarding attributes	91
3.4.32.	Application setup attributes	92
3.4.33.	Attribute support level attributes	92
3.4.34.	Descriptive attributes	92
3.4.35.	Process group attributes	93
3.5.	Callback Functions	93
3.5.1.	Release Callback Function	94
3.5.2.	Modex Callback Function	94
3.5.3.	Spawn Callback Function	95
3.5.4.	Op Callback Function	96
3.5.5.	Lookup Callback Function	96
3.5.6.	Value Callback Function	97
3.5.7.	Info Callback Function	97
3.5.8.	Event Handler Registration Callback Function	98
3.5.9.	Notification Handler Completion Callback Function	99
3.5.10.	Notification Function	100
3.5.11.	Server Setup Application Callback Function	101
3.5.12.	Server Direct Modex Response Callback Function	102
3.5.13.	PMIx Client Connection Callback Function	103
3.5.14.	PMIx Tool Connection Callback Function	103
3.5.15.	Credential callback function	104
3.5.16.	Credential validation callback function	105
3.5.17.	IOF delivery function	106
3.5.18.	IOF and Event registration function	107
3.6.	Constant String Functions	108

4. Initialization and Finalization	111
4.1. Query	111
4.1.1. PMIx_Initialized	111
4.1.2. PMIx_Get_version	112
4.2. Client Initialization and Finalization	112
4.2.1. PMIx_Init	112
4.2.2. PMIx_Finalize	115
4.3. Tool Initialization and Finalization	115
4.3.1. PMIx_tool_init	115
4.3.2. PMIx_tool_finalize	119
4.3.3. PMIx_tool_connect_to_server	119
4.4. Server Initialization and Finalization	121
4.4.1. PMIx_server_init	121
4.4.2. PMIx_server_finalize	123
5. Key/Value Management	124
5.1. Setting and Accessing Key/Value Pairs	124
5.1.1. PMIx_Put	124
5.1.2. PMIx_Get	125
5.1.3. PMIx_Get_nb	128
5.1.4. PMIx_Store_internal	131
5.1.5. Accessing information: examples	132
5.2. Exchanging Key/Value Pairs	136
5.2.1. PMIx_Commit	137
5.2.2. PMIx_Fence	137
5.2.3. PMIx_Fence_nb	139
5.3. Publish and Lookup Data	142
5.3.1. PMIx_Publish	142
5.3.2. PMIx_Publish_nb	144
5.3.3. PMIx_Lookup	145
5.3.4. PMIx_Lookup_nb	148
5.3.5. PMIx_Unpublish	149
5.3.6. PMIx_Unpublish_nb	151

6. Process Management	153
6.1. Abort	153
6.1.1. PMIx_Abort	153
6.2. Process Creation	154
6.2.1. PMIx_Spawn	154
6.2.2. PMIx_Spawn_nb	159
6.3. Connecting and Disconnecting Processes	163
6.3.1. PMIx_Connect	164
6.3.2. PMIx_Connect_nb	166
6.3.3. PMIx_Disconnect	168
6.3.4. PMIx_Disconnect_nb	170
6.4. IO Forwarding	172
6.4.1. PMIx_IOF_pull	173
6.4.2. PMIx_IOF_deregister	175
6.4.3. PMIx_IOF_push	176
7. Job Management and Reporting	179
7.1. Query	179
7.1.1. PMIx_Resolve_peers	179
7.1.2. PMIx_Resolve_nodes	180
7.1.3. PMIx_Query_info	181
7.1.4. PMIx_Query_info_nb	185
7.2. Allocation Requests	191
7.2.1. PMIx_Allocation_request	192
7.2.2. PMIx_Allocation_request_nb	194
7.3. Job Control	197
7.3.1. PMIx_Job_control	198
7.3.2. PMIx_Job_control_nb	200
7.4. Process and Job Monitoring	203
7.4.1. PMIx_Process_monitor	204
7.4.2. PMIx_Process_monitor_nb	206
7.4.3. PMIx_Heartbeat	207
7.5. Logging	208
7.5.1. PMIx_Log	208

7.5.2.	<code>PMIx_Log_nb</code>	210
8.	Event Notification	214
8.1.	Notification and Management	214
8.1.1.	<code>PMIx_Register_event_handler</code>	216
8.1.2.	<code>PMIx_Deregister_event_handler</code>	219
8.1.3.	<code>PMIx_Notify_event</code>	220
9.	Data Packing and Unpacking	223
9.1.	Data Buffer Type	223
9.2.	Support Macros	224
9.2.1.	<code>PMIX_DATA_BUFFER_CREATE</code>	224
9.2.2.	<code>PMIX_DATA_BUFFER_RELEASE</code>	224
9.2.3.	<code>PMIX_DATA_BUFFER_CONSTRUCT</code>	224
9.2.4.	<code>PMIX_DATA_BUFFER_DESTRUCT</code>	225
9.2.5.	<code>PMIX_DATA_BUFFER_LOAD</code>	225
9.2.6.	<code>PMIX_DATA_BUFFER_UNLOAD</code>	226
9.3.	General Routines	226
9.3.1.	<code>PMIx_Data_pack</code>	226
9.3.2.	<code>PMIx_Data_unpack</code>	228
9.3.3.	<code>PMIx_Data_copy</code>	230
9.3.4.	<code>PMIx_Data_print</code>	230
9.3.5.	<code>PMIx_Data_copy_payload</code>	231
10.	Security	233
10.1.	Obtaining Credentials	234
10.1.1.	<code>PMIx_Get_credential</code>	234
10.2.	Validating Credentials	236
10.2.1.	<code>PMIx_Validate_credential</code>	236
11.	Server-Specific Interfaces	238
11.1.	Server Support Functions	238
11.1.1.	<code>PMIx_generate_regex</code>	238
11.1.2.	<code>PMIx_generate_ppn</code>	239
11.1.3.	<code>PMIx_server_register_namespace</code>	240

11.1.4.	<code>PMIx_server_deregister_nspace</code>	253
11.1.5.	<code>PMIx_server_register_client</code>	255
11.1.6.	<code>PMIx_server_deregister_client</code>	256
11.1.7.	<code>PMIx_server_setup_fork</code>	257
11.1.8.	<code>PMIx_server_dmodex_request</code>	257
11.1.9.	<code>PMIx_server_setup_application</code>	259
11.1.10.	<code>PMIx_Register_attributes</code>	261
11.1.11.	<code>PMIx_server_setup_local_support</code>	262
11.1.12.	<code>PMIx_server_IOF_deliver</code>	264
11.1.13.	<code>PMIx_server_collect_inventory</code>	265
11.1.14.	<code>PMIx_server_deliver_inventory</code>	266
11.2.	Server Function Pointers	267
11.2.1.	<code>pmix_server_module_t</code> Module	267
11.2.2.	<code>pmix_server_client_connected_fn_t</code>	268
11.2.3.	<code>pmix_server_client_finalized_fn_t</code>	270
11.2.4.	<code>pmix_server_abort_fn_t</code>	271
11.2.5.	<code>pmix_server_fence_nb_fn_t</code>	272
11.2.6.	<code>pmix_server_dmodex_req_fn_t</code>	276
11.2.7.	<code>pmix_server_publish_fn_t</code>	277
11.2.8.	<code>pmix_server_lookup_fn_t</code>	279
11.2.9.	<code>pmix_server_unpublish_fn_t</code>	281
11.2.10.	<code>pmix_server_spawn_fn_t</code>	283
11.2.11.	<code>pmix_server_connect_fn_t</code>	288
11.2.12.	<code>pmix_server_disconnect_fn_t</code>	290
11.2.13.	<code>pmix_server_register_events_fn_t</code>	291
11.2.14.	<code>pmix_server_deregister_events_fn_t</code>	293
11.2.15.	<code>pmix_server_notify_event_fn_t</code>	295
11.2.16.	<code>pmix_server_listener_fn_t</code>	296
11.2.17.	<code>pmix_server_query_fn_t</code>	297
11.2.18.	<code>pmix_server_tool_connection_fn_t</code>	299
11.2.19.	<code>pmix_server_log_fn_t</code>	300
11.2.20.	<code>pmix_server_alloc_fn_t</code>	302
11.2.21.	<code>pmix_server_job_control_fn_t</code>	305

11.2.22.	<code>pmix_server_monitor_fn_t</code>	308
11.2.23.	<code>pmix_server_get_cred_fn_t</code>	311
11.2.24.	<code>pmix_server_validate_cred_fn_t</code>	312
11.2.25.	<code>pmix_server_iof_fn_t</code>	314
11.2.26.	<code>pmix_server_stdin_fn_t</code>	317
12.	Scheduler-Specific Interfaces	319
12.1.	Scheduler Support Datatypes	319
12.1.1.	Fabric registration structure	319
12.1.2.	Scheduler Support Error Constants	320
12.1.3.	Scheduler Support Attributes	320
12.2.	Scheduler Support Functions	320
12.2.1.	<code>PMIx_server_register_fabric</code>	320
12.2.2.	<code>PMIx_server_deregister_fabric</code>	321
12.2.3.	<code>PMIx_server_get_vertex_info</code>	322
12.2.4.	<code>PMIx_server_get_index</code>	322
13.	Process Sets and Groups	324
13.1.	Process Sets	324
13.2.	Process Groups	325
13.2.1.	<code>PMIx_Group_construct</code>	327
13.2.2.	<code>PMIx_Group_construct_nb</code>	331
13.2.3.	<code>PMIx_Group_destruct</code>	334
13.2.4.	<code>PMIx_Group_destruct_nb</code>	336
13.2.5.	<code>PMIx_Group_invite</code>	338
13.2.6.	<code>PMIx_Group_invite_nb</code>	342
13.2.7.	<code>PMIx_Group_join</code>	344
13.2.8.	<code>PMIx_Group_join_nb</code>	347
13.2.9.	<code>PMIx_Group_leave</code>	349
13.2.10.	<code>PMIx_Group_leave_nb</code>	350
14.	Network Coordinates	352
14.1.	Network Coordinate Datatypes	352
14.1.1.	Network Coordinate Structure	352
14.1.2.	Network Coordinate Support Macros	353

14.1.3.	Network Coordinate Views	354
14.1.4.	Network Coordinate Error Constants	355
14.1.5.	Network Descriptive Attributes	355
A.	Python Bindings	357
A.1.	Datatype Definitions	357
A.1.1.	Example	362
A.2.	Function Definitions	362
A.2.1.	IOF Delivery Function	362
A.2.2.	Event Handler	363
A.2.3.	Server Module Functions	364
A.3.	PMIxClient	375
A.3.1.	Client.init	376
A.3.2.	Client.initialized	376
A.3.3.	Client.get_version	376
A.3.4.	Client.finalize	377
A.3.5.	Client.abort	377
A.3.6.	Client.store_internal	377
A.3.7.	Client.put	378
A.3.8.	Client.commit	378
A.3.9.	Client.fence	379
A.3.10.	Client.get	379
A.3.11.	Client.publish	380
A.3.12.	Client.lookup	380
A.3.13.	Client.unpublish	381
A.3.14.	Client.spawn	381
A.3.15.	Client.connect	382
A.3.16.	Client.disconnect	382
A.3.17.	Client.resolve_peers	383
A.3.18.	Client.resolve_nodes	383
A.3.19.	Client.query	384
A.3.20.	Client.log	384
A.3.21.	Client.allocate	385
A.3.22.	Client.job_ctrl	385

A.3.23.	Client.monitor	386
A.3.24.	Client.get_credential	386
A.3.25.	Client.validate_credential	387
A.3.26.	Client.group_construct	387
A.3.27.	Client.group_invite	388
A.3.28.	Client.group_join	389
A.3.29.	Client.group_leave	389
A.3.30.	Client.group_destruct	390
A.3.31.	Client.register_event_handler	390
A.3.32.	Client.deregister_event_handler	391
A.3.33.	Client.notify_event	391
A.3.34.	Client.error_string	392
A.3.35.	Client.proc_state_string	392
A.3.36.	Client.scope_string	392
A.3.37.	Client.persistence_string	393
A.3.38.	Client.data_range_string	393
A.3.39.	Client.info_directives_string	394
A.3.40.	Client.data_type_string	394
A.3.41.	Client.alloc_directive_string	394
A.3.42.	Client.iof_channel_string	395
A.4.	PMIxServer	395
A.4.1.	Server.init	395
A.4.2.	Server.finalize	396
A.4.3.	Server.generate_regex	396
A.4.4.	Server.generate_ppn	397
A.4.5.	Server.register_nspace	397
A.4.6.	Server.deregister_nspace	398
A.4.7.	Server.register_client	398
A.4.8.	Server.deregister_client	399
A.4.9.	Server.setup_fork	399
A.4.10.	Server.dmodex_request	400
A.4.11.	Server.setup_application	400
A.4.12.	Server.register_attributes	401

A.4.13.	Server.setup_local_support	401
A.4.14.	Server.iof_deliver	402
A.4.15.	Server.collect_inventory	402
A.4.16.	Server.deliver_inventory	403
A.5.	PMIxTool	403
A.5.1.	Tool.init	403
A.5.2.	Tool.finalize	404
A.5.3.	Tool.connect_to_server	404
A.5.4.	Tool.iof_pull	405
A.5.5.	Tool.iof_deregister	405
A.5.6.	Tool.iof_push	406
B.	Acknowledgements	407
B.1.	Version 3.0	407
B.2.	Version 2.0	408
B.3.	Version 1.0	409
	Bibliography	410
	Index	411

CHAPTER 1

Introduction

1 The Process Management Interface (PMI) has been used for quite some time as a means of
2 exchanging wireup information needed for inter-process communication. Two versions (PMI-1 and
3 PMI-2) have been released as part of the MPICH effort, with PMI-2 demonstrating better scaling
4 properties than its PMI-1 predecessor. However, two significant challenges face the High
5 Performance Computing (HPC) community as it continues to move towards machines capable of
6 exaflop and higher performance levels:

- 7 • the physical scale of the machines, and the corresponding number of total processes they support,
8 is expected to reach levels approaching 1 million processes executing across 100 thousand nodes.
9 Prior methods for initiating applications relied on exchanging communication endpoint
10 information between the processes, either directly or in some form of hierarchical collective
11 operation. Regardless of the specific mechanism employed, the exchange across such large
12 applications would consume considerable time, with estimates running in excess of 5-10
13 minutes; and
- 14 • whether it be hybrid applications that combine OpenMP threading operations with MPI, or
15 application-steered workflow computations, the HPC community is experiencing an
16 unprecedented wave of new approaches for computing at exascale levels. One common thread
17 across the proposed methods is an increasing need for orchestration between the application and
18 the system management software stack (SMS) comprising the scheduler (a.k.a. the workload
19 manager (WLM)), the resource manager (RM), global file system, fabric, and other subsystems.
20 The lack of available support for application-to-SMS integration has forced researchers to
21 develop "virtual" environments that hide the SMS behind a customized abstraction layer, but this
22 results in considerable duplication of effort and a lack of portability.

23 Process Management Interface - Exascale (PMIx) represents an attempt to resolve these questions
24 by providing an extended version of the PMI definitions specifically designed to support clusters up
25 to exascale and larger sizes. The overall objective of the project is not to branch the existing
26 definitions – in fact, PMIx fully supports both of the existing PMI-1 and PMI-2 Application
27 Programming Interfaces (APIs) – but rather to:

- 28 a) add flexibility to the existing APIs by adding an array of key-value “attribute” pairs to each API
29 signature that allows implementers to customize the behavior of the API as future needs emerge
30 without having to alter or create new variants of it;
- 31 b) add new APIs that provide extended capabilities such as asynchronous event notification plus
32 dynamic resource allocation and management;

- c) establish a collaboration between SMS subsystem providers including resource manager, fabric, file system, and programming library developers to define integration points between the various subsystems as well as agreed upon definitions for associated APIs, attribute names, and data types;
- d) form a standards-like body for the definitions; and
- e) provide a reference implementation of the PMIx standard.

Complete information about the PMIx standard and affiliated projects can be found at the PMIx web site: <https://pmix.org>

1.1 Charter

The charter of the PMIx community is to:

- Define a set of agnostic APIs (not affiliated with any specific programming model or code base) to support interactions between application processes and the SMS.
- Develop an open source (non-copy-left licensed) standalone “reference” library implementation to facilitate adoption of the PMIx standard.
- Retain transparent backward compatibility with the existing PMI-1 and PMI-2 definitions, any future PMI releases, and across all PMIx versions.
- Support the “Instant On” initiative for rapid startup of applications at exascale and beyond.
- Work with the HPC community to define and implement new APIs that support evolving programming model requirements for application interactions with the SMS.

Participation in the PMIx community is open to anyone, and not restricted to only code contributors to the reference implementation.

1.2 PMIx Standard Overview

The PMIx Standard defines and describes the interface developed by the PMIx Reference Implementation (PRI). Much of this document is specific to the PMIx Reference Implementation (PRI)’s design and implementation. Specifically the standard describes the functionality provided by the PRI, and what the PRI requires of the clients and resource managers (RMs) that use it’s interface.

1.2.1 Who should use the standard?

The PMIx Standard informs PMIx clients and RMs of the syntax and semantics of the PMIx APIs. PMIx clients (e.g., tools, Message Passing Environment (MPE) libraries) can use this standard to understand the set of attributes provided by various APIs of the PRI and their intended behavior.

1 Additional information about the rationale for the selection of specific interfaces and attributes is
2 also provided.

3 PMIx-enabled RMs can use this standard to understand the expected behavior required of them
4 when they support various interfaces/attributes. In addition, optional features and suggestions on
5 behavior are also included in the discussion to help guide RM design and implementation.

6 **1.2.2 What is defined in the standard?**

7 The PMIx Standard defines and describes the interface developed by the PMIx Reference
8 Implementation (PRI). It defines the set of attributes that the PRI supports; the set of attributes that
9 are required of a RM to support, for a given interface; and the set of optional attributes that an RM
10 may choose to support, for a given interface.

11 **1.2.3 What is *not* defined in the standard?**

12 No standards body can require an implementer to support something in their standard, and PMIx is
13 no different in that regard. While an implementer of the PMIx library itself must at least include the
14 standard PMIx headers and instantiate each function, they are free to return “not supported” for any
15 function they choose not to implement.

16 This also applies to the host environments. Resource managers and other system management stack
17 components retain the right to decide on support of a particular function. The PMIx community
18 continues to look at ways to assist SMS implementers in their decisions by highlighting functions
19 that are critical to basic application execution (e.g., [PMIx_Get](#)), while leaving flexibility for
20 tailoring a vendor’s software for their target market segment.

21 One area where this can become more complicated is regarding the attributes that provide
22 information to the client process and/or control the behavior of a PMIx standard API. For example,
23 the [PMIX_TIMEOUT](#) attribute can be used to specify the time (in seconds) before the requested
24 operation should time out. The intent of this attribute is to allow the client to avoid “hanging” in a
25 request that takes longer than the client wishes to wait, or may never return (e.g., a [PMIx_Fence](#)
26 that a blocked participant never enters).

27 If an application (for example) truly relies on the [PMIX_TIMEOUT](#) attribute in a call to
28 [PMIx_Fence](#) , it should set the required flag in the [pmix_info_t](#) for that attribute. This
29 informs the library and its SMS host that it must return an immediate error if this attribute is not
30 supported. By not setting the flag, the library and SMS host are allowed to treat the attribute as
31 optional, ignoring it if support is not available.

32 It is therefore critical that users and application implementers:

- 33 a) consider whether or not a given attribute is required, marking it accordingly; and
- 34 b) check the return status on all PMIx function calls to ensure support was present and that the
35 request was accepted. Note that for non-blocking APIs, a return of [PMIX_SUCCESS](#) only
36 indicates that the request had no obvious errors and is being processed – the eventual callback
37 will return the status of the requested operation itself.

1 While a PMIx library implementer, or an SMS component server, may choose to support a
2 particular PMIx API, they are not required to support every attribute that might apply to it. This
3 would pose a significant barrier to entry for an implementer as there can be a broad range of
4 applicable attributes to a given API, at least some of which may rarely be used. The PMIx
5 community is attempting to help differentiate the attributes by indicating those that are generally
6 used (and therefore, of higher importance to support) vs those that a “complete implementation”
7 would support.

8 Note that an environment that does not include support for a particular attribute/API pair is not
9 “incomplete” or of lower quality than one that does include that support. Vendors must decide
10 where to invest their time based on the needs of their target markets, and it is perfectly reasonable
11 for them to perform cost/benefit decisions when considering what functions and attributes to
12 support.

13 The flip side of that statement is also true: Users who find that their current vendor does not support
14 a function or attribute they require may raise that concern with their vendor and request that the
15 implementation be expanded. Alternatively, users may wish to utilize the PMIx-based Reference
16 RunTime Environment (PRRTE) as a “shim” between their application and the host environment as
17 it might provide the desired support until the vendor can respond. Finally, in the extreme, one can
18 exploit the portability of PMIx-based applications to change vendors.

19 **1.2.4 General Guidance for PMIx Users and Implementors**

20 The PMIx Standard defines the behavior of the PMIx Reference Implementation (PRI). A complete
21 system harnessing the PMIx interface requires an agreement between the PMIx client, be it a tool or
22 library, and the PMIx-enabled RM. The PRI acts as an intermediary between these two entities by
23 providing a standard API for the exchange of requests and responses. The degree to which the
24 PMIx client and the PMIx-enabled RM may interact needs to be defined by those developer
25 communities. The PMIx standard can be used to define the specifics of this interaction.

26 PMIx clients (e.g., tools, MPE libraries) may find that they depend only on a small subset of
27 interfaces and attributes to work correctly. PMIx clients are strongly advised to define a document
28 itemizing the PMIx interfaces and associated attributes that are required for correct operation, and
29 are optional but recommended for full functionality. The PMIx standard cannot define this list for
30 all given PMIx clients, but such a list is valuable to RMs desiring to support these clients.

31 PMIx-enabled RMs may choose to implement a subset of the PMIx standard and/or define attributes
32 beyond those defined herein. PMIx-enabled RMs are strongly advised to define a document
33 itemizing the PMIx interfaces and associated attributes they support, with any annotations about
34 behavior limitations. The PMIx standard cannot define this list for all given PMIx-enabled RMs,
35 but such a list is valuable to PMIx clients desiring to support a broad range of PMIx-enabled RMs.

36 **1.3 PMIx Architecture Overview**

37 This section presents a brief overview of the PMIx Architecture [1]. Note that this is a conceptual
38 model solely used to help guide the standards process — it does not represent a design requirement

1 on any PMIx implementation. Instead, the model is used by the PMIx community as a sounding
 2 board for evaluating proposed interfaces and avoid unintentionally imposing constraints on
 3 implementers. Built into the model are two guiding principles also reflected in the standard. First,
 4 PMIx operates in the mode of a *messenger*, and not a *doer* — i.e., the role of PMIx is to provide
 5 communication between the various participants, relaying requests and returning responses. The
 6 intent of the standard is not to suggest that PMIx itself actually perform any of the defined
 7 operations — this is left to the various SMS elements and/or the application. Any exceptions to that
 8 intent are left to the discretion of the particular implementation.




Figure 1.1.: PMIx-SMS Interactions

9 Thus, as the diagram in Fig. 1.1 shows, the application is built against a PMIx client library that
 10 contains the client-side APIs, attribute definitions, and communication support for interacting with
 11 the local PMIx server. Intra-process cross-library interactions are supported at the client level to
 12 avoid unnecessary burdens on the server. Orchestration requests are sent to the local PMIx server,
 13 which subsequently passes them to the host SMS (here represented by an RM daemon) using the
 14 PMIx server callback functions the host SMS registered during `PMIx_server_init`. The host SMS
 15 can indicate its lack of support for any operation by simply providing a *NULL* for the associated
 16 callback function, or can create a function entry that returns *not supported* when called.

17 The conceptual model places the burden of fulfilling the request on the host SMS. This includes
 18 performing any inter-node communications, or interacting with other SMS elements. Thus, a client
 19 request for a network traffic report does not go directly from the client to the Fabric Manager (FM),
 20 but instead is relayed to the PMIx server, and then passed to the host SMS for execution. This
 21 architecture reflects the second principle underlying the standard — namely, that connectivity is to
 22 be minimized by channeling all application interactions with the SMS through the local PMIx
 23 server.

24 Recognizing the burden this places on SMS vendors, the PMIx community has included interfaces

1 by which the host can request support from local SMS elements. Once the SMS has transferred the
2 request to an appropriate location, a PMIx server interface can be used to pass the request between
3 SMS subsystems. For example, a request for network traffic statistics can utilize the PMIx
4 networking abstractions to retrieve the information from the FM. This reduces the portability and
5 interoperability issues between the individual subsystems by transferring the burden of defining the
6 interoperable interfaces from the SMS subsystems to the PMIx community, which continues to
7 work with those providers to develop the necessary support.

8 Tools, whether standalone or embedded in job scripts, are an exception to the communication rule
9 and can connect to any PMIx server providing they are given adequate rendezvous information. The
10 PMIx conceptual model views the collection of PMIx servers as a cloud-like conglomerate — i.e.,
11 orchestration and information requests can be given to any server regardless of location. However,
12 tools frequently execute on locations that may not house an operating PMIx server — e.g., a users
13 notebook computer. Thus, tools need the ability to remotely connect to the PMIx server “cloud”.

14 The scope of the PMIx standard therefore spans the range of these interactions, between
15 client-and-SMS and between SMS subsystems. Note again that this does not impose a requirement
16 on any given PMIx implementation to cover the entire range — implementers are free to return *not*
17 *supported* from any PMIx function.

18 1.3.1 The PMIx Reference Implementation (PRI)

19 The PMIx community has committed to providing a complete, reference implementation of each
20 version of the standard. Note that the definition of the PMIx Standard is not contingent upon use of
21 the PMIx Reference Implementation (PRI) — any implementation that supports the defined APIs is
22 a PMIx Standard compliant implementation. The PRI is provided solely for the following purposes:

- 23 • Validation of the standard.
24 No proposed change and/or extension to the PMIx standard is accepted without an accompanying
25 prototype implementation in the PRI. This ensures that the proposal has undergone at least some
26 minimal level of scrutiny and testing before being considered.
- 27 • Ease of adoption.
28 The PRI is designed to be particularly easy for resource managers (and the SMS in general) to
29 adopt, thus facilitating a rapid uptake into that community for application portability. Both client
30 and server PMIx libraries are included, along with examples of client usage and server-side
31 integration. A list of supported environments and versions is maintained on the PMIx web site
32 <https://pmix.org/support/faq/what-apis-are-supported-on-my-rm/>

33 The PRI does provide some internal implementations that lie outside the scope of the PMIx
34 standard. This includes several convenience macros as well as support for consolidating collectives
35 for optimization purposes (e.g., the PMIx server aggregates all local **PMIx_Fence** calls before
36 passing them to the SMS for global execution). In a few additional cases, the PMIx community (in
37 partnership with the SMS subsystem providers) have determined that a base level of support for a
38 given operation can best be portably provided by including it in the PRI.

1 Instructions for downloading, and installing the PRI are available on the community’s web site
2 <https://pmix.org/code/getting-the-reference-implementation/>. The PRI targets support for the Linux
3 operating system. A reasonable effort is made to support all major, modern Linux distributions;
4 however, validation is limited to the most recent 2-3 releases of RedHat Enterprise Linux (RHEL),
5 Fedora, CentOS, and SUSE Linux Enterprise Server (SLES). In addition, development support is
6 maintained for Mac OSX. Production support for vendor-specific operating systems is included as
7 provided by the vendor.

8 **1.3.2 The PMIx Reference RunTime Environment (PRRTE)**

9 The PMIx community has also released PRRTE — i.e., a runtime environment containing the
10 reference implementation and capable of operating within a host SMS. PRRTE provides an easy
11 way of exploring PMIx capabilities and testing PMIx-based applications outside of a PMIx-enabled
12 environment by providing a “shim” between the application and the host environment that includes
13 full support for the PRI. The intent of PRRTE is not to replace any existing production
14 environment, but rather to enable developers to work on systems that do not yet feature a
15 PMIx-enabled host SMS or one that lacks a PMIx feature of interest. Instructions for downloading,
16 installing, and using PRRTE are available on the community’s web site
17 <https://pmix.org/code/getting-the-pmix-reference-server/>

18 **1.4 Organization of this document**

19 The remainder of this document is structured as follows:

- 20 • Introduction and Overview in Chapter 1 on page 1
- 21 • Terms and Conventions in Chapter 2 on page 13
- 22 • Data Structures and Types in Chapter 3 on page 19
- 23 • PMIx Initialization and Finalization in Chapter 4 on page 111
- 24 • Key/Value Management in Chapter 5 on page 124
- 25 • Process Management in Chapter 6 on page 153
- 26 • Job Management in Chapter 7 on page 179
- 27 • Event Notification in Chapter 8 on page 214
- 28 • Data Packing and Unpacking in Chapter 9 on page 223
- 29 • PMIx Server Specific Interfaces in Chapter 11 on page 238

1 1.5 Version 1.0: June 12, 2015

2 The PMIx version 1.0 *ad hoc* standard was defined in the PMIx Reference Implementation (PRI)
3 header files as part of the PRI v1.0.0 release prior to the creation of the formal PMIx 2.0 standard.
4 Below are a summary listing of the interfaces defined in the 1.0 headers.

- 5 • Client APIs
 - 6 – `PMIx_Init`, `PMIx_Initialized`, `PMIx_Abort`, `PMIx_Finalize`
 - 7 – `PMIx_Put`, `PMIx_Commit`,
 - 8 – `PMIx_Fence`, `PMIx_Fence_nb`
 - 9 – `PMIx_Get`, `PMIx_Get_nb`
 - 10 – `PMIx_Publish`, `PMIx_Publish_nb`
 - 11 – `PMIx_Lookup`, `PMIx_Lookup`
 - 12 – `PMIx_Unpublish`, `PMIx_Unpublish_nb`
 - 13 – `PMIx_Spawn`, `PMIx_Spawn_nb`
 - 14 – `PMIx_Connect`, `PMIx_Connect_nb`
 - 15 – `PMIx_Disconnect`, `PMIx_Disconnect_nb`
 - 16 – `PMIx_Resolve_nodes`, `PMIx_Resolve_peers`
- 17 • Server APIs
 - 18 – `PMIx_server_init`, `PMIx_server_finalize`
 - 19 – `PMIx_generate_regex`, `PMIx_generate_ppn`
 - 20 – `PMIx_server_register_namespace`, `PMIx_server_deregister_namespace`
 - 21 – `PMIx_server_register_client`, `PMIx_server_deregister_client`
 - 22 – `PMIx_server_setup_fork`, `PMIx_server_dmodex_request`
- 23 • Common APIs
 - 24 – `PMIx_Get_version`, `PMIx_Store_internal`, `PMIx_Error_string`
 - 25 – `PMIx_Register_errhandler`, `PMIx_Deregister_errhandler`, `PMIx_Notify_error`

26 The `PMIx_Init` API was subsequently modified in the PRI release v1.1.0.

27 1.6 Version 2.0: Sept. 2018

28 The following APIs were introduced in v2.0 of the PMIx Standard:

- 1 • Client APIs
- 2 – `PMIx_Query_info_nb` , `PMIx_Log_nb`
- 3 – `PMIx_Allocation_request_nb` , `PMIx_Job_control_nb` ,
- 4 `PMIx_Process_monitor_nb` , `PMIx_Heartbeat`
- 5 • Server APIs
- 6 – `PMIx_server_setup_application` , `PMIx_server_setup_local_support`
- 7 • Tool APIs
- 8 – `PMIx_tool_init` , `PMIx_tool_finalize`
- 9 • Common APIs
- 10 – `PMIx_Register_event_handler` , `PMIx_Deregister_event_handler`
- 11 – `PMIx_Notify_event`
- 12 – `PMIx_Proc_state_string` , `PMIx_Scope_string`
- 13 – `PMIx_Persistence_string` , `PMIx_Data_range_string`
- 14 – `PMIx_Info_directives_string` , `PMIx_Data_type_string`
- 15 – `PMIx_Alloc_directive_string`
- 16 – `PMIx_Data_pack` , `PMIx_Data_unpack` , `PMIx_Data_copy`
- 17 – `PMIx_Data_print` , `PMIx_Data_copy_payload`

18 The `PMIx_Init` API was modified in v2.0 of the standard from its *ad hoc* v1.0 signature to
 19 include passing of a `pmix_info_t` array for flexibility and “future-proofing” of the API. In
 20 addition, the `PMIx_Notify_error`, `PMIx_Register_errhandler`, and `PMIx_Deregister_errhandler`
 21 APIs were replaced.

22 1.7 Version 2.1: Dec. 2018

23 The v2.1 update includes clarifications and corrections from the v2.0 document, plus addition of
 24 examples:

- 25 • Clarify description of `PMIx_Connect` and `PMIx_Disconnect` APIs.
- 26 • Explain that values for the `PMIX_COLLECTIVE_ALGO` are environment-dependent
- 27 • Identify the namespace/rank values required for retrieving attribute-associated information using
 28 the `PMIx_Get` API
- 29 • Provide definitions for `session` , `job` , `application` , and other terms used throughout the
 30 document

- 1 • Clarify definitions of `PMIX_UNIV_SIZE` versus `PMIX_JOB_SIZE`
- 2 • Clarify server module function return values
- 3 • Provide examples of the use of `PMIx_Get` for retrieval of information
- 4 • Clarify the use of `PMIx_Get` versus `PMIx_Query_info_nb`
- 5 • Clarify return values for non-blocking APIs and emphasize that callback functions must not be
- 6 invoked prior to return from the API
- 7 • Provide detailed example for construction of the `PMIx_server_register_nspace` input
- 8 information array
- 9 • Define information levels (e.g., `session` vs `job`) and associated attributes for both storing
- 10 and retrieving values
- 11 • Clarify roles of PMIx server library and host environment for collective operations
- 12 • Clarify definition of `PMIX_UNIV_SIZE`

13 1.8 Version 2.2: Jan 2019

14 The v2.2 update includes the following clarifications and corrections from the v2.1 document:

- 15 • Direct modex upcall function (`pmix_server_dmodex_req_fn_t`) cannot complete
- 16 atomically as the API cannot return the requested information except via the provided callback
- 17 function
- 18 • Add missing `pmix_data_array_t` definition and support macros
- 19 • Add a rule divider between implementer and host environment required attributes for clarity
- 20 • Add `PMIX_QUERY_QUALIFIERS_CREATE` macro to simplify creation of `pmix_query_t`
- 21 qualifiers
- 22 • Add `PMIX_APP_INFO_CREATE` macro to simplify creation of `pmix_app_t` directives
- 23 • Add flag and `PMIX_INFO_IS_END` macro for marking and detecting the end of a
- 24 `pmix_info_t` array
- 25 • Clarify the allowed hierarchical nesting of the `PMIX_SESSION_INFO_ARRAY` ,
- 26 `PMIX_JOB_INFO_ARRAY` , and associated attributes

27 1.9 Version 3.0: Dec. 2018

28 The following APIs were introduced in v3.0 of the PMIx Standard:

- 29 • Client APIs
- 30 – `PMIx_Log` , `PMIx_Job_control`

- 1 – `PMIx_Allocation_request` , `PMIx_Process_monitor`
- 2 – `PMIx_Get_credential` , `PMIx_Validate_credential`
- 3 • Server APIs
- 4 – `PMIx_server_IOF_deliver`
- 5 – `PMIx_server_collect_inventory` , `PMIx_server_deliver_inventory`
- 6 • Tool APIs
- 7 – `PMIx_IOF_pull` , `PMIx_IOF_push` , `PMIx_IOF_deregister`
- 8 – `PMIx_tool_connect_to_server`
- 9 • Common APIs
- 10 – `PMIx_IOF_channel_string`

11 The document added a chapter on security credentials, a new section for Input/Output (IO)
 12 forwarding to the Process Management chapter, and a few blocking forms of previously-existing
 13 non-blocking APIs. Attributes supporting the new APIs were introduced, as well as additional
 14 attributes for a few existing functions.

15 1.10 Version 3.1: Jan. 2019

16 The v3.1 update includes clarifications and corrections from the v3.0 document:

- 17 • Direct modex upcall function (`pmix_server_dmodex_req_fn_t`) cannot complete
 18 atomically as the API cannot return the requested information except via the provided callback
 19 function
- 20 • Fix typo in name of `PMIX_FWD_STDDIAG` attribute
- 21 • Correctly identify the information retrieval and storage attributes as “new” to v3 of the standard
- 22 • Add missing `pmix_data_array_t` definition and support macros
- 23 • Add a rule divider between implementer and host environment required attributes for clarity
- 24 • Add `PMIX_QUERY_QUALIFIERS_CREATE` macro to simplify creation of `pmix_query_t`
 25 qualifiers
- 26 • Add `PMIX_APP_INFO_CREATE` macro to simplify creation of `pmix_app_t` directives
- 27 • Add new attributes to specify the level of information being requested where ambiguity may exist
 28 (see 3.4.11)
- 29 • Add new attributes to assemble information by its level for storage where ambiguity may exist
 30 (see 3.4.12)

- 1 • Add flag and `PMIX_INFO_IS_END` macro for marking and detecting the end of a
2 `pmix_info_t` array
- 3 • Clarify that `PMIX_NUM_SLOTS` is duplicative of (a) `PMIX_UNIV_SIZE` when used at the
4 `session` level and (b) `PMIX_MAX_PROCS` when used at the `job` and `application`
5 levels, but leave it in for backward compatibility.
- 6 • Clarify difference between `PMIX_JOB_SIZE` and `PMIX_MAX_PROCS`
- 7 • Clarify that `PMIx_server_setup_application` must be called per- `job` instead of per-
8 `application` as the name implies. Unfortunately, this is a historical artifact. Note that both
9 `PMIX_NODE_MAP` and `PMIX_PROC_MAP` must be included as input in the `info` array provided
10 to that function. Further descriptive explanation of the “instant on” procedure will be provided in
11 the next version of the PMIx Standard.
- 12 • Clarify how the PMIx server expects data passed to the host by
13 `pmix_server_fence_nb_fn_t` should be aggregated across nodes, and provide a code
14 snippet example

15 1.11 Version 4.0: June 2019

16 The following changes were introduced in v4.0 of the PMIx Standard:

- 17 • Clarified that the `PMIx_Fence_nb` operation can immediately return
18 `PMIX_OPERATION_SUCCEEDED` in lieu of passing the request to a PMIx server if only the
19 calling process is involved in the operation
- 20 • Added the `PMIx_Register_attributes` API by which a host environment can register
21 the attributes it supports for each server-to-host operation
- 22 • Added the ability to query supported attributes from the PMIx tool, client and server libraries, as
23 well as the host environment via the new `pmix_regattr_t` structure. Both human-readable
24 and machine-parsable output is supported. New attributes to support this operation include:
 - 25 – `PMIX_CLIENT_ATTRIBUTES` , `PMIX_SERVER_ATTRIBUTES` ,
26 `PMIX_TOOL_ATTRIBUTES` , and `PMIX_HOST_ATTRIBUTES` to identify which library
27 supports the attribute; and
 - 28 – `PMIX_MAX_VALUE` , `PMIX_MIN_VALUE` , and `PMIX_ENUM_VALUE` to provide
29 machine-parsable description of accepted values

CHAPTER 2

PMIx Terms and Conventions

1 The PMIx Standard has adopted the widespread use of key-value *attributes* to add flexibility to the
2 functionality expressed in the existing APIs. Accordingly, the community has chosen to require that
3 the definition of each standard API include the passing of an array of attributes. These provide a
4 means of customizing the behavior of the API as future needs emerge without having to alter or
5 create new variants of it. In addition, attributes provide a mechanism by which researchers can
6 easily explore new approaches to a given operation without having to modify the API itself.

7 The PMIx community has further adopted a policy that modification of existing released APIs will
8 only be permitted under extreme circumstances. In its effort to avoid introduction of any such
9 backward incompatibility, the community has avoided the definitions of large numbers of APIs that
10 each focus on a narrow scope of functionality, and instead relied on the definition of fewer generic
11 APIs that include arrays of directives for “tuning” the function’s behavior. Thus, modifications to
12 the PMIx standard increasingly consist of the definition of new attributes along with a description
13 of the APIs to which they relate and the expected behavior when used with those APIs.

14 One area where this can become more complicated relates to the attributes that provide directives to
15 the client process and/or control the behavior of a PMIx standard API. For example, the
16 **PMIX_TIMEOUT** attribute can be used to specify the time (in seconds) before the requested
17 operation should time out. The intent of this attribute is to allow the client to avoid hanging in a
18 request that takes longer than the client wishes to wait, or may never return (e.g., a **PMIx_Fence**
19 that a blocked participant never enters).

20 If an application truly relies on the **PMIX_TIMEOUT** attribute in a call to **PMIx_Fence** , it
21 should set the *required* flag in the **pmix_info_t** for that attribute. This informs the library and
22 its SMS host that it must return an immediate error if this attribute is not supported. By not setting
23 the flag, the library and SMS host are allowed to treat the attribute as optional, silently ignoring it if
24 support is not available.

Advice to users

25 It is critical that users and application developers consider whether or not a given attribute is
26 required (marking it accordingly) and always check the return status on all PMIx function calls to
27 ensure support was present and that the request was accepted. Note that for non-blocking APIs, a
28 return of **PMIX_SUCCESS** only indicates that the request had no obvious errors and is being
29 processed. The eventual callback will return the status of the requested operation itself.

1 While a PMIx library implementer, or an SMS component server, may choose to support a
2 particular PMIx API, they are not required to support every attribute that might apply to it. This
3 would pose a significant barrier to entry for an implementer as there can be a broad range of
4 applicable attributes to a given API, at least some of which may rarely be used in a specific market
5 area. The PMIx community is attempting to help differentiate the attributes by indicating in the
6 standard those that are generally used (and therefore, of higher importance to support) versus those
7 that a “complete implementation” would support.

8 In addition, the document refers to the following entities and process stages when describing
9 use-cases or operations involving PMIx:

- 10 • *session* refers to an allocated set of resources assigned to a particular user by the system WLM.
11 Historically, HPC sessions have consisted of a static allocation of resources - i.e., a block of
12 resources are assigned to a user in response to a specific request and managed as a unified
13 collection. However, this is changing in response to the growing use of dynamic programming
14 models that require on-the-fly allocation and release of system resources. Accordingly, the term
15 *session* in this document refers to the current block of assigned resources and is a potentially
16 dynamic entity.
- 17 • *slot* refers to an allocated entry for a process. WLMs frequently allocate entire nodes to a
18 *session*, but can also be configured to define the maximum number of processes that can
19 simultaneously be executed on each node. This often corresponds to the number of hardware
20 Processing Units (PUs) (typically cores, but can also be defined as hardware threads) on the
21 node. However, the correlation between hardware PUs and slot allocations strictly depends upon
22 system configuration.
- 23 • *job* refers to a set of one or more *applications* executed as a single invocation by the user within a
24 session. For example, “*mpiexec -n 1 app1 : -n 2 app2*” is considered a single Multiple Program
25 Multiple Data (MPMD) job containing two applications.
- 26 • *namespace* refers to a character string value assigned by the RM to a *job*. All *applications*
27 executed as part of that *job* share the same *namespace*. The *namespace* assigned to each *job* must
28 be unique within the scope of the governing RM.
- 29 • *application* refers to a single executable (binary, script, etc.) member of a *job*. Applications
30 consist of one or more *processes*, either operating independently or in parallel at any given time
31 during their execution.
- 32 • *rank* refers to the numerical location (starting from zero) of a process within the defined scope.
33 Thus, global rank is the rank of a process within its *job*, while *application rank* is the rank of that
34 process within its *application*.
- 35 • *workflow* refers to an orchestrated execution plan frequently spanning multiple *jobs* carried out
36 under the control of a *workflow manager* process. An example workflow might first execute a
37 computational job to generate the flow of liquid through a complex cavity, followed by a
38 visualization job that takes the output of the first job as its input to produce an image output.

- 1 • *scheduler* refers to the component of the SMS responsible for scheduling of resource allocations.
2 This is also generally referred to as the *system workflow manager* - for the purposes of this
3 document, the *WLM* acronym will be used interchangeably to refer to the scheduler.
- 4 • *resource manager* is used in a generic sense to represent the subsystem that will host the PMIx
5 server library. This could be a vendor's RM, a programming library's RunTime
6 Environment (RTE), or some other agent.
- 7 • *host environment* is used interchangeably with *resource manager* to refer to the process hosting
8 the PMIx server library.
- 9 • *network plane* refers to a collection of Network Interface Cards (NICs) and switches in a
10 common logical or physical configuration. Network planes are often implemented in HPC
11 clusters as separate overlay or physical networks controlled by a dedicated fabric manager.

12 This document borrows freely from other standards (most notably from the Message Passing
13 Interface (MPI) and OpenMP standards) in its use of notation and conventions in an attempt to
14 reduce confusion. The following sections provide an overview of the conventions used throughout
15 the PMIx Standard document.

16 2.1 Notational Conventions

17 Some sections of this document describe programming language specific examples or APIs. Text
18 that applies only to programs for which the base language is C is shown as follows:


19  C

C specific text...

20 `int foo = 42;`

21  C

22 Some text is for information only, and is not part of the normative specification. These take several
23 forms, described in their examples below:

24 

Note: General text...

25 

26  Rationale

24 Throughout this document, the rationale for the design choices made in the interface specification is
25 set off in this section. Some readers may wish to skip these sections, while readers interested in
26 interface design may want to read them carefully.



Advice to users

1 Throughout this document, material aimed at users and that illustrates usage is set off in this
2 section. Some readers may wish to skip these sections, while readers interested in programming
3 with the PMIx API may want to read them carefully.

Advice to PMIx library implementers

4 Throughout this document, material that is primarily commentary to PMIx library implementers is
5 set off in this section. Some readers may wish to skip these sections, while readers interested in
6 PMIx implementations may want to read them carefully.

Advice to PMIx server hosts

7 Throughout this document, material that is primarily commentary aimed at host environments (e.g.,
8 RMs and RTEs) providing support for the PMIx server library is set off in this section. Some
9 readers may wish to skip these sections, while readers interested in integrating PMIx servers into
10 their environment may want to read them carefully.

2.2 Semantics

12 The following terms will be taken to mean:

- 13 • *shall*, *must* and *will* indicate that the specified behavior is *required* of all conforming
14 implementations
- 15 • *should* and *may* indicate behaviors that a complete implementation would include, but are not
16 required of all conforming implementations

1 2.3 Naming Conventions

2 The PMIx standard has adopted the following conventions:

- 3 • PMIx constants and attributes are prefixed with **PMIX**.
- 4 • Structures and type definitions are prefixed with **pmix**.
- 5 • Underscores are used to separate words in a function or variable name.
- 6 • Lowercase letters are used in PMIx client APIs except for the PMIx prefix (noted below) and the
7 first letter of the word following it. For example, **PMIx_Get_version** .
- 8 • PMIx server and tool APIs are all lower case letters following the prefix - e.g.,
9 **PMIx_server_register_namespace** .
- 10 • The **PMIx_** prefix is used to denote functions.
- 11 • The **pmix_** prefix is used to denote function pointer and type definitions.

12 Users should not use the **PMIX**, **PMIx**, or **pmix** prefixes in their applications or libraries so as to
13 avoid symbol conflicts with current and later versions of the PMIx standard and implementations
14 such as the PRI.

15 2.4 Procedure Conventions

16 While the current PMIx Reference Implementation (PRI) is solely based on the C programming
17 language, it is not the intent of the PMIx Standard to preclude the use of other languages.
18 Accordingly, the procedure specifications in the PMIx Standard are written in a
19 language-independent syntax with the arguments marked as IN, OUT, or INOUT. The meanings of
20 these are:

- 21 • IN: The call may use the input value but does not update the argument from the perspective of
22 the caller at any time during the calls execution,
- 23 • OUT: The call may update the argument but does not use its input value
- 24 • INOUT: The call may both use and update the argument.

25 Many PMIx interfaces, particularly nonblocking interfaces, use a **void*cbdata** object passed to
26 the function that is then passed to the associated callback. In a client-side API, the cbdata is a
27 client-provided context (opaque object) that the client can pass to the nonblocking call (e.g.,
28 **PMIx_Get_nb**). When the nonblocking call (e.g., **pmix_value_cbfunc_t**) completes, the
29 cbdata is passed back to the client without modification by the PMIx library, thus allowing the
30 client to associate a context with that callback. This is useful if there are many outstanding
31 nonblocking calls.

32 A similar model is used for the server module functions (see 11.2.1). In this case, the PMIx library
33 is making an upcall into its host via the PMIx server module function and passing a specific cbfunc

1 and cbdata. The PMIx library expects the host to call the cbfunc with the necessary arguments and
2 pass back the original cbdata upon completing the operation. This gives the server-side PMIx
3 library the ability to associate a context with the call back (since multiple operations may be
4 outstanding). The host has no visibility into the contents of the cbdata object, nor is permitted to
5 alter it in any way.

6 2.5 Standard vs Reference Implementation

7 The *PMIx Standard* is implementation independent. The *PMIx Reference Implementation* (PRI) is
8 one implementation of the Standard and the PMIx community strives to ensure that it fully
9 implements the Standard. Given its role as the community's testbed and its widespread use, this
10 document cites the attributes supported by the PRI for each API where relevant by marking them in
11 red. This is not meant to imply nor confer any special role to the PRI with respect to the Standard
12 itself, but instead to provide a convenience to users of the Standard and PRI.

13 Similarly, the *PMIx Reference RunTime Environment* (PRRTE) is provided by the community to
14 enable users operating in non-PMIx environments to develop and execute PMIx-enabled
15 applications and tools. Attributes supported by the PRRTE are marked in green.

CHAPTER 3

Data Structures and Types

1 This chapter defines PMIx standard data structures (along with macros for convenient use), types,
2 and constants. These apply to all consumers of the PMIx interface. Where necessary for
3 clarification, the description of, for example, an attribute may be copied from this chapter into a
4 section where it is used.

5 A PMIx implementation may define additional attributes beyond those specified in this document.

▼ Advice to PMIx library implementers ▼

6 Structures, types, and macros in the PMIx Standard are defined in terms of the C-programming
7 language. Implementers wishing to support other languages should provide the equivalent
8 definitions in a language-appropriate manner.

9 If a PMIx implementation chooses to define additional attributes they should avoid using the **PMIX**
10 prefix in their name or starting the attribute string with a *pmix* prefix. This helps the end user
11 distinguish between what is defined by the PMIx standard and what is specific to that PMIx
12 implementation, and avoids potential conflicts with attributes defined by the standard.



▼ Advice to users ▼

13 Use of increment/decrement operations on indices inside PMIx macros is discouraged due to
14 unpredictable behavior. For example, the following sequence:

```
15 PMIX_INFO_LOAD(&array[n++], "mykey", &mystring, PMIX_STRING);  
16 PMIX_INFO_LOAD(&array[n++], "mykey2", &myint, PMIX_INT);
```

17 will load the given key-values into incorrect locations if the macro is implemented as:

```
18 define PMIX_INFO_LOAD(m, k, v, t) \\  
19     do { \\  
20         if (NULL != (k)) { \\  
21             pmix_strncpy((m)->key, (k), PMIX_MAX_KEYLEN); \\  
22         } \\  
23         (m)->flags = 0; \\  
24         pmix_value_load(&((m)->value), (v), (t)); \\  
25     } while (0)
```

26 since the index is cited more than once in the macro. The PMIx standard only governs the existence
27 and syntax of macros - it does not specify their implementation. Given the freedom of
28 implementation, a safer call sequence might be as follows:

```
1     PMIX_INFO_LOAD(&array[n], "mykey", &mystring, PMIX_STRING);
2     ++n;
3     PMIX_INFO_LOAD(&array[n], "mykey2", &myint, PMIX_INT);
4     ++n;
```

5 3.1 Constants

6 PMIx defines a few values that are used throughout the standard to set the size of fixed arrays or as
7 a means of identifying values with special meaning. The community makes every attempt to
8 minimize the number of such definitions. The constants defined in this section may be used before
9 calling any PMIx library initialization routine. Additional constants associated with specific data
10 structures or types are defined in the section describing that data structure or type.

11 **PMIX_MAX_NSLEN** Maximum namespace string length as an integer.

▼ Advice to PMIx library implementers ▼

12 **PMIX_MAX_NSLEN** should have a minimum value of 63 characters. Namespace arrays in PMIx
13 defined structures must reserve a space of size **PMIX_MAX_NSLEN** +1 to allow room for the **NULL**
14 terminator

15 **PMIX_MAX_KEYLEN** Maximum key string length as an integer.

▼ Advice to PMIx library implementers ▼

16 **PMIX_MAX_KEYLEN** should have a minimum value of 63 characters. Key arrays in PMIx defined
17 structures must reserve a space of size **PMIX_MAX_KEYLEN** +1 to allow room for the **NULL**
18 terminator

1 3.1.1 PMIx Error Constants

2 The `pmix_status_t` structure is an `int` type for return status.

3 The tables shown in this section define the possible values for `pmix_status_t`. PMIx errors are
4 required to always be negative, with 0 reserved for `PMIX_SUCCESS`. Values in the list that were
5 deprecated in later standards are denoted as such. Values added to the list in this version of the
6 standard are shown in **magenta**.

▼ Advice to PMIx library implementers ▼

7 A PMIx implementation must define all of the constants defined in this section, even if they will
8 never return the specific value to the caller.



▼ Advice to users ▼

9 Other than `PMIX_SUCCESS` (which is required to be zero), the actual value of any PMIx error
10 constant is left to the PMIx library implementer. Thus, users are advised to always refer to constant
11 by name, and not a specific implementation's value, for portability between implementations and
12 compatibility across library versions.



13 3.1.1.1 General Error Constants

14 These are general constants originally defined in versions 1 and 2 of the PMIx Standard.

15	<code>PMIX_SUCCESS</code>	Success
16	<code>PMIX_ERROR</code>	General Error
17	<code>PMIX_ERR_SILENT</code>	Silent error
18	<code>PMIX_ERR_DEBUGGER_RELEASE</code>	Error in debugger release
19	<code>PMIX_ERR_PROC_RESTART</code>	Fault tolerance: Error in process restart
20	<code>PMIX_ERR_PROC_CHECKPOINT</code>	Fault tolerance: Error in process checkpoint
21	<code>PMIX_ERR_PROC_MIGRATE</code>	Fault tolerance: Error in process migration
22	<code>PMIX_ERR_PROC_ABORTED</code>	Process was aborted
23	<code>PMIX_ERR_PROC_REQUESTED_ABORT</code>	Process is already requested to abort
24	<code>PMIX_ERR_PROC_ABORTING</code>	Process is being aborted
25	<code>PMIX_ERR_SERVER_FAILED_REQUEST</code>	Failed to connect to the server
26	<code>PMIX_EXISTS</code>	Requested operation would overwrite an existing value
27	<code>PMIX_ERR_INVALID_CRED</code>	Invalid security credentials
28	<code>PMIX_ERR_HANDSHAKE_FAILED</code>	Connection handshake failed
29	<code>PMIX_ERR_READY_FOR_HANDSHAKE</code>	Ready for handshake
30	<code>PMIX_ERR_WOULD_BLOCK</code>	Operation would block
31	<code>PMIX_ERR_UNKNOWN_DATA_TYPE</code>	Unknown data type
32	<code>PMIX_ERR_PROC_ENTRY_NOT_FOUND</code>	Process not found
33	<code>PMIX_ERR_TYPE_MISMATCH</code>	Invalid type
34	<code>PMIX_ERR_UNPACK_INADEQUATE_SPACE</code>	Inadequate space to unpack data

1 **PMIX_ERR_UNPACK_FAILURE** Unpack failed
2 **PMIX_ERR_PACK_FAILURE** Pack failed
3 **PMIX_ERR_PACK_MISMATCH** Pack mismatch
4 **PMIX_ERR_NO_PERMISSIONS** No permissions
5 **PMIX_ERR_TIMEOUT** Timeout expired
6 **PMIX_ERR_UNREACH** Unreachable
7 **PMIX_ERR_IN_ERRNO** Error defined in **errno**
8 **PMIX_ERR_BAD_PARAM** Bad parameter
9 **PMIX_ERR_RESOURCE_BUSY** Resource busy
10 **PMIX_ERR_OUT_OF_RESOURCE** Resource exhausted
11 **PMIX_ERR_DATA_VALUE_NOT_FOUND** Data value not found
12 **PMIX_ERR_INIT** Error during initialization
13 **PMIX_ERR_NOMEM** Out of memory
14 **PMIX_ERR_INVALID_ARG** Invalid argument
15 **PMIX_ERR_INVALID_KEY** Invalid key
16 **PMIX_ERR_INVALID_KEY_LENGTH** Invalid key length
17 **PMIX_ERR_INVALID_VAL** Invalid value
18 **PMIX_ERR_INVALID_VAL_LENGTH** Invalid value length
19 **PMIX_ERR_INVALID_LENGTH** Invalid argument length
20 **PMIX_ERR_INVALID_NUM_ARGS** Invalid number of arguments
21 **PMIX_ERR_INVALID_ARGS** Invalid arguments
22 **PMIX_ERR_INVALID_NUM_PARSED** Invalid number parsed
23 **PMIX_ERR_INVALID_KEYVALP** Invalid key/value pair
24 **PMIX_ERR_INVALID_SIZE** Invalid size
25 **PMIX_ERR_INVALID_NAMESPACE** Invalid namespace
26 **PMIX_ERR_SERVER_NOT_AVAIL** Server is not available
27 **PMIX_ERR_NOT_FOUND** Not found
28 **PMIX_ERR_NOT_SUPPORTED** Not supported
29 **PMIX_ERR_NOT_IMPLEMENTED** Not implemented
30 **PMIX_ERR_COMM_FAILURE** Communication failure
31 **PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER** Unpacking past the end of the buffer
32 provided
33 **PMIX_ERR_LOST_CONNECTION_TO_SERVER** Lost connection to server
34 **PMIX_ERR_LOST_PEER_CONNECTION** Lost connection to peer
35 **PMIX_ERR_LOST_CONNECTION_TO_CLIENT** Lost connection to client
36 **PMIX_QUERY_PARTIAL_SUCCESS** Query partial success (used by query system)
37 **PMIX_NOTIFY_ALLOC_COMPLETE** Notify that allocation is complete
38 **PMIX_JCTRL_CHECKPOINT** Job control: Monitored by PMIx client to trigger checkpoint
39 operation
40 **PMIX_JCTRL_CHECKPOINT_COMPLETE** Job control: Sent by PMIx client and monitored
41 by PMIx server to notify that requested checkpoint operation has completed.
42 **PMIX_JCTRL_PREEMPT_ALERT** Job control: Monitored by PMIx client to detect an RM
43 intending to preempt the job.

1 **PMIX_MONITOR_HEARTBEAT_ALERT** Job monitoring: Heartbeat alert
 2 **PMIX_MONITOR_FILE_ALERT** Job monitoring: File alert
 3 **PMIX_PROC_TERMINATED** Process terminated - can be either normal or abnormal
 4 termination
 5 **PMIX_ERR_INVALID_TERMINATION** Process terminated without calling
 6 **PMIx_Finalize**, or was a member of an assemblage formed via **PMIx_Connect** and
 7 terminated or called **PMIx_Finalize** without first calling **PMIx_Disconnect** (or its
 8 non-blocking form) from that assemblage.

9 3.1.1.2 Operational Error Constants

10 **PMIX_ERR_EVENT_REGISTRATION** Error in event registration
 11 **PMIX_ERR_JOB_TERMINATED** Error job terminated
 12 **PMIX_ERR_UPDATE_ENDPOINTS** Error updating endpoints
 13 **PMIX_MODEL_DECLARED** Model declared
 14 **PMIX_GDS_ACTION_COMPLETE** The global data storage (GDS) action has completed
 15 **PMIX_ERR_INVALID_OPERATION** The requested operation is supported by the
 16 implementation and host environment, but fails to meet a requirement (e.g., requesting to
 17 *disconnect* from processes without first *connecting* to them).
 18 **PMIX_PROC_HAS_CONNECTED** A tool or client has connected to the PMIx server
 19 **PMIX_CONNECT_REQUESTED** Connection has been requested by a PMIx-based tool
 20 **PMIX_MODEL_RESOURCES** Resource usage by a programming model has changed
 21 **PMIX_OPENMP_PARALLEL_ENTERED** An OpenMP parallel code region has been entered
 22 **PMIX_OPENMP_PARALLEL_EXITED** An OpenMP parallel code region has completed
 23 **PMIX_LAUNCH_DIRECTIVE** Launcher directives have been received from a PMIx-enabled
 24 tool
 25 **PMIX_LAUNCHER_READY** Application launcher (e.g., mpiexec) is ready to receive directives
 26 from a PMIx-enabled tool
 27 **PMIX_LAUNCH_COMPLETE** A job has been launched - the nspace of the launched job will be
 28 included in the notification
 29 **PMIX_OPERATION_IN_PROGRESS** A requested operation is already in progress
 30 **PMIX_OPERATION_SUCCEEDED** The requested operation was performed atomically - no
 31 callback function will be executed
 32 **PMIX_ERR_PARTIAL_SUCCESS** The operation is considered successful but not all elements
 33 of the operation were concluded (e.g., some members of a group construct operation chose
 34 not to participate)
 35 **PMIX_ERR_DUPLICATE_KEY** The provided key has already been published on a different
 36 data range
 37 **PMIX_ERR_INVALID_OPERATION** The requested operation is not valid - this can possibly
 38 indicate the inclusion of conflicting directives or a request to perform an operation that
 39 conflicts with an ongoing one.
 40 **PMIX_GROUP_INVITED** The process has been invited to join a PMIx Group - the identifier of
 41 the group and the ID's of other invited (or already joined) members will be included in the
 42 notification

1 **PMIX_GROUP_LEFT** A process has asynchronously left a PMIx Group - the process identifier
2 of the departing process will be included in the notification

3 **PMIX_GROUP_MEMBER_FAILED** A member of a PMIx Group has abnormally terminated
4 (i.e., without formally leaving the group prior to termination) - the process identifier of the
5 failed process will be included in the notification

6 **PMIX_GROUP_INVITE_ACCEPTED** A process has accepted an invitation to join a PMIx
7 Group - the identifier of the group being joined will be included in the notification

8 **PMIX_GROUP_INVITE_DECLINED** A process has declined an invitation to join a PMIx
9 Group - the identifier of the declined group will be included in the notification

10 **PMIX_GROUP_INVITE_FAILED** An invited process failed or terminated prior to responding
11 to the invitation - the identifier of the failed process will be included in the notification.

12 **PMIX_GROUP_MEMBERSHIP_UPDATE** The membership of a PMIx group has changed - the
13 identifiers of the revised membership will be included in the notification.

14 **PMIX_GROUP_CONSTRUCT_ABORT** Any participant in a PMIx group construct operation
15 that returns **PMIX_GROUP_CONSTRUCT_ABORT** from the *leader failed* event handler will
16 cause all participants to receive an event notifying them of that status. Similarly, the leader
17 may elect to abort the procedure by either returning this error code from the handler assigned
18 to the **PMIX_GROUP_INVITE_ACCEPTED** or **PMIX_GROUP_INVITE_DECLINED**
19 codes, or by generating an event for the abort code. Abort events will be sent to all invited or
20 existing members of the group.

21 **PMIX_GROUP_CONSTRUCT_COMPLETE** The group construct operation has completed - the
22 final membership will be included in the notification.

23 **PMIX_GROUP_LEADER_FAILED** The current *leader* of a group including this process has
24 abnormally terminated - the group identifier will be included in the notification.

25 **PMIX_GROUP_LEADER_SELECTED** A new *leader* of a group including this process has been
26 selected - the identifier of the new leader will be included in the notification

27 **PMIX_GROUP_CONTEXT_ID_ASSIGNED** A new Process Group Context
28 Identifier (PGCID) has been assigned by the host environment to a group that includes this
29 process - the group identifier will be included in the notification.

30 **PMIX_ERR_REPEAT_ATTR_REGISTRATION** The attributes for an identical function have
31 already been registered at the specified level (host, server, or client)

32 **PMIX_ERR_IOF_FAILURE** An IO forwarding operation failed - the affected channel will be
33 included in the notification

34 **PMIX_ERR_IOF_COMPLETE** IO forwarding of the standard input for this process has
35 completed - i.e., the stdin file descriptor has closed

36 3.1.1.3 System error constants

37 **PMIX_ERR_SYS_BASE** Mark the beginning of a dedicated range of constants for system event
38 reporting.

39 **PMIX_ERR_NODE_DOWN** A node has gone down - the identifier of the affected node will be
40 included in the notification

41 **PMIX_ERR_NODE_OFFLINE** A node has been marked as *offline* - the identifier of the affected
42 node will be included in the notification

1 **PMIX_ERR_SYS_OTHER** Mark the end of a dedicated range of constants for system event
2 reporting.

3 3.1.1.4 Event handler error constants

4 **PMIX_EVENT_NO_ACTION_TAKEN** Event handler: No action taken
5 **PMIX_EVENT_PARTIAL_ACTION_TAKEN** Event handler: Partial action taken
6 **PMIX_EVENT_ACTION_DEFERRED** Event handler: Action deferred
7 **PMIX_EVENT_ACTION_COMPLETE** Event handler: Action complete

8 3.1.1.5 User-Defined Error Constants

9 PMIx establishes an error code boundary for constants defined in the PMIx standard. Negative
10 values larger than this (and any positive values greater than zero) are guaranteed not to conflict with
11 PMIx values.


12 **PMIX_EXTERNAL_ERR_BASE** A starting point for user-level defined error constants.
13 Negative values lower than this are guaranteed not to conflict with PMIx values. Definitions
14 should always be based on the **PMIX_EXTERNAL_ERR_BASE** constant and not a specific
15 value as the value of the constant may change.

16 3.1.2 Macros for use with PMIx constants

17 3.1.2.1 Detect system event constant

18 Test a given error constant to see if it falls within the dedicated range of constants for system event
19 reporting.

PMIx v2.2



21 **IN a**
22 Error constant to be checked (**pmix_status_t**)

23 Returns **true** if the provided values falls within the dedicated range of constants for system event
24 reporting

25 3.2 Data Types

26 This section defines various data types used by the PMIx APIs. The version of the standard in
27 which a particular data type was introduced is shown in the margin.

1 3.2.1 Key Structure

2 The `pmix_key_t` structure is a statically defined character array of length `PMIX_MAX_KEYLEN`
3 +1, thus supporting keys of maximum length `PMIX_MAX_KEYLEN` while preserving space for a
4 mandatory `NULL` terminator.

PMIx v2.0

```
▼ typedef char pmix_key_t [PMIX_MAX_KEYLEN+1]; ▲  
C
```

6 Characters in the key must be standard alphanumeric values supported by common utilities such as
7 *strcmp*.

▼ **Advice to users** ▲

8 References to keys in PMIx v1 were defined simply as an array of characters of size
9 `PMIX_MAX_KEYLEN+1`. The `pmix_key_t` type definition was introduced in version 2 of the
10 standard. The two definitions are code-compatible and thus do not represent a break in backward
11 compatibility.

12 Passing a `pmix_key_t` value to the standard *sizeof* utility can result in compiler warnings of
13 incorrect returned value. Users are advised to avoid using *sizeof(pmix_key_t)* and instead rely on
14 the `PMIX_MAX_KEYLEN` constant.

▲

15 3.2.1.1 Key support macro

16 Compare the key in a `pmix_info_t` to a given value

PMIx v3.0

```
▼ PMIX_CHECK_KEY(a, b) ▲  
C
```

18 **IN a**
19 Pointer to the structure whose key is to be checked (pointer to `pmix_info_t`)

20 **IN b**
21 String value to be compared against (`char*`)

22 Returns `true` if the key matches the given value

1 3.2.2 Namespace Structure

2 The `pmix_namespace_t` structure is a statically defined character array of length
3 `PMIX_MAX_NSLEN + 1`, thus supporting namespaces of maximum length `PMIX_MAX_NSLEN`
4 while preserving space for a mandatory `NULL` terminator.

PMIx v2.0

```
▼ C ▼  
5 typedef char pmix_namespace_t [PMIX_MAX_NSLEN+1];  
▲ C ▲
```

6 Characters in the namespace must be standard alphanumeric values supported by common utilities
7 such as `strcmp`.

▼ Advice to users ▼

8 References to namespace values in *PMIx v1* were defined simply as an array of characters of size
9 `PMIX_MAX_NSLEN+1`. The `pmix_namespace_t` type definition was introduced in version 2 of the
10 standard. The two definitions are code-compatible and thus do not represent a break in backward
11 compatibility.

12 Passing a `pmix_namespace_t` value to the standard `sizeof` utility can result in compiler warnings of
13 incorrect returned value. Users are advised to avoid using `sizeof(pmix_namespace_t)` and instead rely
14 on the `PMIX_MAX_NSLEN` constant.

▲

15 3.2.2.1 Namespace support macro

16 Compare the string in a `pmix_namespace_t` to a given value

PMIx v3.0

```
▼ C ▼  
17 PMIX_CHECK_NAMESPACE (a, b)  
▲ C ▲
```

18 **IN a**
19 Pointer to the structure whose value is to be checked (pointer to `pmix_namespace_t`)

20 **IN b**
21 String value to be compared against (`char*`)

22 Returns `true` if the namespace matches the given value

1 3.2.3 Rank Structure

2 The `pmix_rank_t` structure is a `uint32_t` type for rank values.

PMIx v1.0

```
3 typedef uint32_t pmix_rank_t;
```

4 The following constants can be used to set a variable of the type `pmix_rank_t`. All definitions
5 were introduced in version 1 of the standard unless otherwise marked. Valid rank values start at
6 zero.

7 **PMIX_RANK_UNDEF** A value to request job-level data where the information itself is not
8 associated with any specific rank, or when passing a `pmix_proc_t` identifier to an
9 operation that only references the namespace field of that structure.

10 **PMIX_RANK_WILDCARD** A value to indicate that the user wants the data for the given key
11 from every rank that posted that key.

12 **PMIX_RANK_LOCAL_NODE** Special rank value used to define groups of ranks. This constant
13 defines the group of all ranks on a local node.

14 **PMIX_RANK_LOCAL_PEERS** Special rank value used to define groups of rankss. This
15 constant defines the group of all ranks on a local node within the same namespace as the
16 current process.

17 **PMIX_RANK_INVALID** An invalid rank value.

18 **PMIX_RANK_VALID** Define an upper boundary for valid rank values.

19 3.2.4 Process Structure

20 The `pmix_proc_t` structure is used to identify a single process in the PMIx universe. It contains
21 a reference to the namespace and the `pmix_rank_t` within that namespace.

PMIx v1.0

```
22 typedef struct pmix_proc {  
23     pmix_nspace_t nspace;  
24     pmix_rank_t rank;  
25 } pmix_proc_t;
```

26 3.2.5 Process structure support macros

27 The following macros are provided to support the `pmix_proc_t` structure.

1 **3.2.5.1 Initialize the `pmix_proc_t` structure**

2 `PMIX_PROC_CONSTRUCT`

3 Initialize the `pmix_proc_t` fields

PMIx v1.0



4 `PMIX_PROC_CONSTRUCT (m)`



5 **IN** `m`

6 Pointer to the structure to be initialized (pointer to `pmix_proc_t`)

7 **3.2.5.2 Destruct the `pmix_proc_t` structure**

8 There is nothing to release here as the fields in `pmix_proc_t` are all declared *static*. However,
9 the macro is provided for symmetry in the code and for future-proofing should some allocated field
10 be included some day.

11 **3.2.5.3 Create a `pmix_proc_t` array**

12 Allocate and initialize an array of `pmix_proc_t` structures

PMIx v1.0



13 `PMIX_PROC_CREATE (m, n)`



14 **INOUT** `m`

15 Address where the pointer to the array of `pmix_proc_t` structures shall be stored (handle)

16 **IN** `n`

17 Number of structures to be allocated (`size_t`)

18 **3.2.5.4 Free a `pmix_proc_t` array**

19 Release an array of `pmix_proc_t` structures

PMIx v1.0



20 `PMIX_PROC_FREE (m, n)`



21 **IN** `m`

22 Pointer to the array of `pmix_proc_t` structures (handle)

23 **IN** `n`

24 Number of structures in the array (`size_t`)

1 3.2.5.5 Load a `pmix_proc_t` structure

2 Load values into a `pmix_proc_t`

PMIx v2.0



C

3 `PMIX_PROC_LOAD (m, n, r)`



C

4 **IN** `m`

Pointer to the structure to be loaded (pointer to `pmix_proc_t`)

6 **IN** `n`

Namespace to be loaded (`pmix_namespace_t`)

8 **IN** `r`

Rank to be assigned (`pmix_rank_t`)

10 3.2.5.6 Compare identifiers

11 Compare two `pmix_proc_t` identifiers

PMIx v3.0



C

12 `PMIX_CHECK_PROCID (a, b)`



C

13 **IN** `a`

Pointer to a structure whose ID is to be compared (pointer to `pmix_proc_t`)

15 **IN** `b`

Pointer to a structure whose ID is to be compared (pointer to `pmix_proc_t`)

17 Returns `true` if the two structures contain matching namespaces and:

- 18 • the ranks are the same value
- 19 • one of the ranks is `PMIX_RANK_WILDCARD`

20 3.2.6 Process State Structure

21 *PMIx v2.0* The `pmix_proc_state_t` structure is a `uint8_t` type for process state values. The following
22 constants can be used to set a variable of the type `pmix_proc_state_t`. All values were
23 originally defined in version 2 of the standard unless otherwise marked.



Advice to users

24 The fine-grained nature of the following constants may exceed the ability of an RM to provide
25 updated process state values during the process lifetime. This is particularly true of states in the
26 launch process, and for short-lived processes.

1 **PMIX_PROC_STATE_UNDEF** Undefined process state
2 **PMIX_PROC_STATE_PREPPED** Process is ready to be launched
3 **PMIX_PROC_STATE_LAUNCH_UNDERWAY** Process launch is underway
4 **PMIX_PROC_STATE_RESTART** Process is ready for restart
5 **PMIX_PROC_STATE_TERMINATE** Process is marked for termination
6 **PMIX_PROC_STATE_RUNNING** Process has been locally **fork**'ed by the RM
7 **PMIX_PROC_STATE_CONNECTED** Process has connected to PMIx server
8 **PMIX_PROC_STATE_UNTERMINATED** Define a "boundary" between the terminated states
9 and **PMIX_PROC_STATE_CONNECTED** so users can easily and quickly determine if a
10 process is still running or not. Any value less than this constant means that the process has not
11 terminated.
12 **PMIX_PROC_STATE_TERMINATED** Process has terminated and is no longer running
13 **PMIX_PROC_STATE_ERROR** Define a boundary so users can easily and quickly determine if
14 a process abnormally terminated. Any value above this constant means that the process has
15 terminated abnormally.
16 **PMIX_PROC_STATE_KILLED_BY_CMD** Process was killed by a command
17 **PMIX_PROC_STATE_ABORTED** Process was aborted by a call to **PMIx_Abort**
18 **PMIX_PROC_STATE_FAILED_TO_START** Process failed to start
19 **PMIX_PROC_STATE_ABORTED_BY_SIG** Process aborted by a signal
20 **PMIX_PROC_STATE_TERM_WO_SYNC** Process exited without calling **PMIx_Finalize**
21 **PMIX_PROC_STATE_COMM_FAILED** Process communication has failed
22 **PMIX_PROC_STATE_CALLED_ABORT** Process called **PMIx_Abort**
23 **PMIX_PROC_STATE_MIGRATING** Process failed and is waiting for resources before
24 restarting
25 **PMIX_PROC_STATE_CANNOT_RESTART** Process failed and cannot be restarted
26 **PMIX_PROC_STATE_TERM_NON_ZERO** Process exited with a non-zero status
27 **PMIX_PROC_STATE_FAILED_TO_LAUNCH** Unable to launch process

28 3.2.7 Process Information Structure

29 The **pmix_proc_info_t** structure defines a set of information about a specific process
30 including its name, location, and state.

PMIx v2.0

```

1  typedef struct pmix_proc_info {
2      /** Process structure */
3      pmix_proc_t proc;
4      /** Hostname where process resides */
5      char *hostname;
6      /** Name of the executable */
7      char *executable_name;
8      /** Process ID on the host */
9      pid_t pid;
10     /** Exit code of the process. Default: 0 */
11     int exit_code;
12     /** Current state of the process */
13     pmix_proc_state_t state;
14 } pmix_proc_info_t;

```

3.2.8 Process Information Structure support macros

The following macros are provided to support the `pmix_proc_info_t` structure.

3.2.8.1 Initialize the `pmix_proc_info_t` structure

Initialize the `pmix_proc_info_t` fields

PMIx v2.0

```

19  PMIX_PROC_INFO_CONSTRUCT(m)

```

IN `m`
 Pointer to the structure to be initialized (pointer to `pmix_proc_info_t`)

3.2.8.2 Destruct the `pmix_proc_info_t` structure

Destruct the `pmix_proc_info_t` fields

PMIx v2.0

```

24  PMIX_PROC_INFO_DESTRUCT(m)

```

IN `m`
 Pointer to the structure to be destructed (pointer to `pmix_proc_info_t`)

1 3.2.8.3 Create a `pmix_proc_info_t` array

2 Allocate and initialize a `pmix_proc_info_t` array

PMIx v2.0

▼ `C` ▼

3 `PMIX_PROC_INFO_CREATE (m, n)`

▲ `C` ▲

4 **INOUT** `m`

5 Address where the pointer to the array of `pmix_proc_info_t` structures shall be stored
6 (handle)

7 **IN** `n`

8 Number of structures to be allocated (`size_t`)

9 3.2.8.4 Free a `pmix_proc_info_t` array

10 Release an array of `pmix_proc_info_t` structures

PMIx v2.0

▼ `C` ▼

11 `PMIX_PROC_INFO_FREE (m, n)`

▲ `C` ▲

12 **IN** `m`

13 Pointer to the array of `pmix_proc_info_t` structures (handle)

14 **IN** `n`

15 Number of structures in the array (`size_t`)

16 3.2.9 Scope of Put Data

17 *PMIx v1.0* The `pmix_scope_t` structure is a `uint8_t` type that defines the scope for data passed to
18 `PMIx_Put`. The following constants can be used to set a variable of the type `pmix_scope_t`.
19 All definitions were introduced in version 1 of the standard unless otherwise marked.

20 Specific implementations may support different scope values, but all implementations must support
21 at least `PMIX_GLOBAL`. If a scope value is not supported, then the `PMIx_Put` call must return
22 `PMIX_ERR_NOT_SUPPORTED`.

23 **PMIX_SCOPE_UNDEF** Undefined scope

24 **PMIX_LOCAL** The data is intended only for other application processes on the same node.

25 Data marked in this way will not be included in data packages sent to remote requestors —
26 i.e., it is only available to processes on the local node.

27 **PMIX_REMOTE** The data is intended solely for applications processes on remote nodes. Data
28 marked in this way will not be shared with other processes on the same node — i.e., it is only
29 available to processes on remote nodes.

30 **PMIX_GLOBAL** The data is to be shared with all other requesting processes, regardless of
31 location.

32 *PMIx v2.0* **PMIX_INTERNAL** The data is intended solely for this process and is not shared with other
33 processes.

1 3.2.10 Range of Published Data

2 *PMIx v1.0* The `pmix_data_range_t` structure is a `uint8_t` type that defines a range for data *published*
3 via functions other than `PMIx_Put` - e.g., the `PMIx_Publish` API. The following constants
4 can be used to set a variable of the type `pmix_data_range_t`. Several values were initially
5 defined in version 1 of the standard but subsequently renamed and other values added in version 2.
6 Thus, all values shown below are as they were defined in version 2 except where noted.

7 `PMIX_RANGE_UNDEF` Undefined range
8 `PMIX_RANGE_RM` Data is intended for the host resource manager.
9 `PMIX_RANGE_LOCAL` Data is only available to processes on the local node.
10 `PMIX_RANGE_NAMESPACE` Data is only available to processes in the same namespace.
11 `PMIX_RANGE_SESSION` Data is only available to all processes in the session.
12 `PMIX_RANGE_GLOBAL` Data is available to all processes.
13 `PMIX_RANGE_CUSTOM` Range is specified in the `pmix_info_t` associated with this call.
14 `PMIX_RANGE_PROC_LOCAL` Data is only available to this process.
15 `PMIX_RANGE_INVALID` Invalid value

Advice to users

16 The names of the `pmix_data_range_t` values changed between version 1 and version 2 of the
17 standard, thereby breaking backward compatibility

18 3.2.11 Data Persistence Structure

19 *PMIx v1.0* The `pmix_persistence_t` structure is a `uint8_t` type that defines the policy for data
20 published by clients via the `PMIx_Publish` API. The following constants can be used to set a
21 variable of the type `pmix_persistence_t`. All definitions were introduced in version 1 of the
22 standard unless otherwise marked.

23 `PMIX_PERSIST_INDEF` Retain data until specifically deleted.
24 `PMIX_PERSIST_FIRST_READ` Retain data until the first access, then the data is deleted.
25 `PMIX_PERSIST_PROC` Retain data until the publishing process terminates.
26 `PMIX_PERSIST_APP` Retain data until the application terminates.
27 `PMIX_PERSIST_SESSION` Retain data until the session/allocation terminates.
28 `PMIX_PERSIST_INVALID` Invalid value

1 3.2.12 Data Array Structure

PMIx v2.0

```
2 typedef struct pmix_data_array
3     pmix_data_type_t type;
4     size_t size;
5     void *array;
6     pmix_data_array_t;
```

C

C

7 The `pmix_data_array_t` structure is used to pass arrays of related values. Any PMIx data
8 type (including complex structures) can be included in the array.

9 3.2.13 Data array structure support macros

10 The following macros are provided to support the `pmix_data_array_t` structure.

11 3.2.13.1 Initialize the `pmix_data_array_t` structure

12 Initialize the `pmix_data_array_t` fields, allocating memory for the array itself.

PMIx v2.2

```
13 PMIX_DATA_ARRAY_CONSTRUCT(m, n, t)
```

C

C

```
14 IN m
15     Pointer to the structure to be initialized (pointer to pmix_data_array_t)
16 IN n
17     Number of elements in the array (size_t)
18 IN t
19     PMIx data type for the array elements (pmix_data_type_t)
```

20 3.2.13.2 Destruct the `pmix_data_array_t` structure

21 Destruct the `pmix_data_array_t` fields, releasing the array's memory.

PMIx v2.2

```
22 PMIX_DATA_ARRAY_DESTRUCT(m)
```

C

C

```
23 IN m
24     Pointer to the structure to be destructed (pointer to pmix_data_array_t)
```

1 3.2.13.3 Create and initialize a `pmix_data_array_t` object

2 Allocate and initialize a `pmix_data_array_t` structure and initialize it, allocating memory for
3 the array itself as well.

PMIx v2.2

▼ `C` _____ ▼
4 **PMIX_DATA_ARRAY_CREATE**(*m*, *n*, *t*)
▲ _____ ▲ `C`

5 **INOUT** *m*

6 Address where the pointer to the `pmix_data_array_t` structure shall be stored (handle)

7 **IN** *n*

8 Number of elements in the array (`size_t`)

9 **IN** *t*

10 PMIx data type for the array elements (`pmix_data_type_t`)

11 3.2.13.4 Free a `pmix_data_array_t` object

12 Release a `pmix_data_array_t` structure, including releasing the array's memory.

PMIx v2.2

▼ `C` _____ ▼
13 **PMIX_DATA_ARRAY_FREE**(*m*)
▲ _____ ▲ `C`

14 **IN** *m*

15 Pointer to the `pmix_data_array_t` structure (handle)

16 3.2.14 Value Structure

17 The `pmix_value_t` structure is used to represent the value passed to `PMIx_Put` and retrieved
18 by `PMIx_Get` , as well as many of the other PMIx functions.

19 A collection of values may be specified under a single key by passing a `pmix_value_t`
20 containing an array of type `pmix_data_array_t` , with each array element containing its own
21 object. All members shown below were introduced in version 1 of the standard unless otherwise
22 marked.

PMIx v1.0

```

1  typedef struct pmix_value {
2      pmix_data_type_t type;
3      union {
4          bool flag;
5          uint8_t byte;
6          char *string;
7          size_t size;
8          pid_t pid;
9          int integer;
10         int8_t int8;
11         int16_t int16;
12         int32_t int32;
13         int64_t int64;
14         unsigned int uint;
15         uint8_t uint8;
16         uint16_t uint16;
17         uint32_t uint32;
18         uint64_t uint64;
19         float fval;
20         double dval;
21         struct timeval tv;
22         time_t time; // version 2.0
23         pmix_status_t status; // version 2.0
24         pmix_rank_t rank; // version 2.0
25         pmix_proc_t *proc; // version 2.0
26         pmix_byte_object_t bo;
27         pmix_persistence_t persist; // version 2.0
28         pmix_scope_t scope; // version 2.0
29         pmix_data_range_t range; // version 2.0
30         pmix_proc_state_t state; // version 2.0
31         pmix_proc_info_t *pinfo; // version 2.0
32         pmix_data_array_t *darray; // version 2.0
33         void *ptr; // version 2.0
34         pmix_alloc_directive_t adir; // version 2.0
35     } data;
36 } pmix_value_t;

```

37 3.2.15 Value structure support macros

38 The following macros are provided to support the `pmix_value_t` structure.

1 3.2.15.1 Initialize the `pmix_value_t` structure

2 Initialize the `pmix_value_t` fields

PMIx v1.0

▼  C 

3 **PMIX_VALUE_CONSTRUCT** (m)

▲  C 

4 **IN** m

5 Pointer to the structure to be initialized (pointer to `pmix_value_t`)

6 3.2.15.2 Destruct the `pmix_value_t` structure

7 Destruct the `pmix_value_t` fields

PMIx v1.0

▼  C 

8 **PMIX_VALUE_DESTRUCT** (m)

▲  C 

9 **IN** m

10 Pointer to the structure to be destructed (pointer to `pmix_value_t`)

11 3.2.15.3 Create a `pmix_value_t` array

12 Allocate and initialize an array of `pmix_value_t` structures

PMIx v1.0

▼  C 

13 **PMIX_VALUE_CREATE** (m, n)

▲  C 

14 **INOUT** m

15 Address where the pointer to the array of `pmix_value_t` structures shall be stored (handle)

16 **IN** n

17 Number of structures to be allocated (`size_t`)

18 3.2.15.4 Free a `pmix_value_t` array

19 Release an array of `pmix_value_t` structures

PMIx v1.0

▼  C 

20 **PMIX_VALUE_FREE** (m, n)

▲  C 

21 **IN** m

22 Pointer to the array of `pmix_value_t` structures (handle)

23 **IN** n

24 Number of structures in the array (`size_t`)

1 3.2.15.5 Load a value structure

2 Summary

3 Load data into a `pmix_value_t` structure.

PMIx v2.0

```
4 PMIX_VALUE_LOAD(v, d, t);
```

5 **IN** v

6 The `pmix_value_t` into which the data is to be loaded (pointer to `pmix_value_t`)

7 **IN** d

8 Pointer to the data value to be loaded (handle)

9 **IN** t

10 Type of the provided data value (`pmix_data_type_t`)

11 Description

12 This macro simplifies the loading of data into a `pmix_value_t` by correctly assigning values to
13 the structure's fields.

Advice to users

14 The data will be copied into the `pmix_value_t` - thus, any data stored in the source value can be
15 modified or free'd without affecting the copied data once the macro has completed.

16 3.2.15.6 Unload a `pmix_value_t` structure

17 Summary

18 Unload data from a `pmix_value_t` structure.

PMIx v2.2

```
19 PMIX_VALUE_UNLOAD(r, v, d, t);
```

20 **OUT** r

21 Status code indicating result of the operation `pmix_status_t`

22 **IN** v

23 The `pmix_value_t` from which the data is to be unloaded (pointer to `pmix_value_t`)

24 **INOUT** d

25 Pointer to the location where the data value is to be returned (handle)

26 **INOUT** t

27 Pointer to return the data type of the unloaded value (handle)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

Description

This macro simplifies the unloading of data from a `pmix_value_t`.

▼ Advice to users ▼

Memory will be allocated and the data will be in the `pmix_value_t` returned - the source `pmix_value_t` will not be altered.

▲

3.2.15.7 Transfer data between `pmix_value_t` structures

Summary

Transfer the data value between two `pmix_value_t` structures.

PMIx v2.0

▼ C ▼

```
PMIX_VALUE_XFER(r, d, s);
```

▲ C ▲

OUT `r`

Status code indicating success or failure of the transfer (`pmix_status_t`)

IN `d`

Pointer to the `pmix_value_t` destination (handle)

IN `s`

Pointer to the `pmix_value_t` source (handle)

Description

This macro simplifies the transfer of data between two `pmix_value_t` structures, ensuring that all fields are properly copied.

▼ Advice to users ▼

The data will be copied into the destination `pmix_value_t` - thus, any data stored in the source value can be modified or free'd without affecting the copied data once the macro has completed.

▲

1 3.2.15.8 Retrieve a numerical value from a `pmix_value_t`

2 Retrieve a numerical value from a `pmix_value_t` structure

PMIx v3.0

```
3 PMIX_VALUE_GET_NUMBER(s, m, n, t)
```

4 **OUT** s

Status code for the request (`pmix_status_t`)

6 **IN** m

Pointer to the `pmix_value_t` structure (handle)

8 **OUT** n

Variable to be set to the value (match expected type)

10 **IN** t

Type of number expected in m (`pmix_data_type_t`)

12 Sets the provided variable equal to the numerical value contained in the given `pmix_value_t`,
13 returning success if the data type of the value matches the expected type and

14 `PMIX_ERR_BAD_PARAM` if it doesn't

15 3.2.16 Info Structure

16 The `pmix_info_t` structure defines a key/value pair with associated directive. All fields were
17 defined in version 1.0 unless otherwise marked.

PMIx v1.0

```
18 typedef struct pmix_info_t {  
19     pmix_key_t key;  
20     pmix_info_directives_t flags;    // version 2.0  
21     pmix_value_t value;  
22 } pmix_info_t;
```

23 3.2.17 Info structure support macros

24 The following macros are provided to support the `pmix_info_t` structure.

25 3.2.17.1 Initialize the `pmix_info_t` structure

26 Initialize the `pmix_info_t` fields

PMIx v1.0

```
27 PMIX_INFO_CONSTRUCT(m)
```

28 **IN** m

29 Pointer to the structure to be initialized (pointer to `pmix_info_t`)

1 **3.2.17.2 Destruct the `pmix_info_t` structure**

2 Destruct the `pmix_info_t` fields

PMIx v1.0



C

3 **PMIX_INFO_DESTRUCT (m)**



C

4 **IN m**

5 Pointer to the structure to be destructed (pointer to `pmix_info_t`)

6 **3.2.17.3 Create a `pmix_info_t` array**

7 Allocate and initialize an array of `pmix_info_t` structures

PMIx v1.0



C

8 **PMIX_INFO_CREATE (m, n)**



C

9 **INOUT m**

10 Address where the pointer to the array of `pmix_info_t` structures shall be stored (handle)

11 **IN n**

12 Number of structures to be allocated (`size_t`)

13 **3.2.17.4 Free a `pmix_info_t` array**

14 Release an array of `pmix_info_t` structures

PMIx v1.0



C

15 **PMIX_INFO_FREE (m, n)**



C

16 **IN m**

17 Pointer to the array of `pmix_info_t` structures (handle)

18 **IN n**

19 Number of structures in the array (`size_t`)

1 3.2.17.5 Load key and value data into a `pmix_info_t`

PMIx v1.0

```
2 PMIX_INFO_LOAD(v, k, d, t);
```

3 **IN** v

4 Pointer to the `pmix_info_t` into which the key and data are to be loaded (pointer to
5 `pmix_info_t`)

6 **IN** k

7 String key to be loaded - must be less than or equal to `PMIX_MAX_KEYLEN` in length
8 (handle)

9 **IN** d

10 Pointer to the data value to be loaded (handle)

11 **IN** t

12 Type of the provided data value (`pmix_data_type_t`)

13 This macro simplifies the loading of key and data into a `pmix_info_t` by correctly assigning
14 values to the structure's fields.

Advice to users

15 Both key and data will be copied into the `pmix_info_t` - thus, the key and any data stored in the
16 source value can be modified or free'd without affecting the copied data once the macro has
17 completed.

18 3.2.17.6 Copy data between `pmix_info_t` structures

19 Copy all data (including key, value, and directives) between two `pmix_info_t` structures.

PMIx v2.0

```
20 PMIX_INFO_XFER(d, s);
```

21 **IN** d

22 Pointer to the destination `pmix_info_t` (pointer to `pmix_info_t`)

23 **IN** s

24 Pointer to the source `pmix_info_t` (pointer to `pmix_info_t`)

25 This macro simplifies the transfer of data between two `pmix_info_t` structures.

Advice to users

26 All data (including key, value, and directives) will be copied into the destination `pmix_info_t` -
27 thus, the source `pmix_info_t` may be free'd without affecting the copied data once the macro
28 has completed.

1 3.2.17.7 Test a boolean `pmix_info_t`

2 A special macro for checking if a boolean `pmix_info_t` is `true`

PMIx v2.0

▼ `C` ▼

3 `PMIX_INFO_TRUE(m)`

▲ `C` ▲

4 **IN** `m`

5 Pointer to a `pmix_info_t` structure (handle)

6 A `pmix_info_t` structure is considered to be of type `PMIX_BOOL` and value `true` if:

- 7
- the structure reports a type of `PMIX_UNDEF`, or
 - the structure reports a type of `PMIX_BOOL` and the data flag is `true`
- 8

9 3.2.18 Info Type Directives

10 *PMIx v2.0* The `pmix_info_directives_t` structure is a `uint32_t` type that defines the behavior of
11 command directives via `pmix_info_t` arrays. By default, the values in the `pmix_info_t`
12 array passed to a PMIx are *optional*.

▼ **Advice to users** ▼

13 A PMIx implementation or PMIx-enabled RM may ignore any `pmix_info_t` value passed to a
14 PMIx API if it is not explicitly marked as `PMIX_INFO_REQD`. This is because the values
15 specified default to optional, meaning they can be ignored. This may lead to unexpected behavior if
16 the user is relying on the behavior specified by the `pmix_info_t` value. If the user relies on the
17 behavior defined by the `pmix_info_t` then they must set the `PMIX_INFO_REQD` flag using the
18 `PMIX_INFO_REQUIRED` macro.

▲

▼ **Advice to PMIx library implementers** ▼

19 The top 16-bits of the `pmix_info_directives_t` are reserved for internal use by PMIx
20 library implementers - the PMIx standard will *not* specify their intent, leaving them for customized
21 use by implementers. Implementers are advised to use the provided `PMIX_INFO_IS_REQUIRED`
22 macro for testing this flag, and must return `PMIX_ERR_NOT_SUPPORTED` as soon as possible to
23 the caller if the required behavior is not supported.

1 The following constants were introduced in version 2.0 (unless otherwise marked) and can be used
2 to set a variable of the type `pmix_info_directives_t`.

3 **PMIX_INFO_REQD** The behavior defined in the `pmix_info_t` array is required, and not
4 optional. This is a bit-mask value.

5 **PMIX_INFO_ARRAY_END** Mark that this `pmix_info_t` struct is at the end of an array
6 created by the `PMIX_INFO_CREATE` macro. This is a bit-mask value.

Advice to PMIx server hosts

7 Host environments are advised to use the provided `PMIX_INFO_IS_REQUIRED` macro for
8 testing this flag and must return `PMIX_ERR_NOT_SUPPORTED` as soon as possible to the caller
9 if the required behavior is not supported.

10 3.2.19 Info Directive support macros

11 The following macros are provided to support the setting and testing of `pmix_info_t` directives.

12 3.2.19.1 Mark an info structure as required

13 Summary

14 Set the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure.

PMIx v2.0

```
15 PMIX_INFO_REQUIRED(info);
```

16 **IN** `info`

17 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

18 This macro simplifies the setting of the `PMIX_INFO_REQD` flag in `pmix_info_t` structures.

19 3.2.19.2 Mark an info structure as optional

20 Summary

21 Unsets the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure.

PMIx v2.0

```
22 PMIX_INFO_OPTIONAL(info);
```

23 **IN** `info`

24 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

25 This macro simplifies marking a `pmix_info_t` structure as *optional*.

1 3.2.19.3 Test an info structure for *required* directive

2 Summary

3 Test the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure, returning `true` if the flag is set.

PMIx v2.0

```
▼ _____ C _____ ▼  
4 PMIX_INFO_IS_REQUIRED(info);  
▲ _____ C _____ ▲
```

5 **IN** `info`

6 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

7 This macro simplifies the testing of the required flag in `pmix_info_t` structures.

8 3.2.19.4 Test an info structure for *optional* directive

9 Summary

10 Test a `pmix_info_t` structure, returning `true` if the structure is *optional*.

PMIx v2.0

```
▼ _____ C _____ ▼  
11 PMIX_INFO_IS_OPTIONAL(info);  
▲ _____ C _____ ▲
```

12 **IN** `info`

13 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

14 Test the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure, returning `true` if the flag is *not*
15 set.

16 3.2.19.5 Test an info structure for *end of array* directive

17 Summary

18 Test a `pmix_info_t` structure, returning `true` if the structure is at the end of an array created
19 by the `PMIX_INFO_CREATE` macro.

PMIx v2.2

```
▼ _____ C _____ ▼  
20 PMIX_INFO_IS_END(info);  
▲ _____ C _____ ▲
```

21 **IN** `info`

22 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

23 This macro simplifies the testing of the end-of-array flag in `pmix_info_t` structures.

1 3.2.20 Job Allocation Directives

2 *PMIx v2.0* The `pmix_alloc_directive_t` structure is a `uint8_t` type that defines the behavior of
3 allocation requests. The following constants can be used to set a variable of the type
4 `pmix_alloc_directive_t`. All definitions were introduced in version 2 of the standard
5 unless otherwise marked.

6 **PMIX_ALLOC_NEW** A new allocation is being requested. The resulting allocation will be
7 disjoint (i.e., not connected in a job sense) from the requesting allocation.

8 **PMIX_ALLOC_EXTEND** Extend the existing allocation, either in time or as additional
9 resources.

10 **PMIX_ALLOC_RELEASE** Release part of the existing allocation. Attributes in the
11 accompanying `pmix_info_t` array may be used to specify permanent release of the
12 identified resources, or “lending” of those resources for some period of time.

13 **PMIX_ALLOC_REACQUIRE** Reacquire resources that were previously “lent” back to the
14 scheduler.

15 **PMIX_ALLOC_EXTERNAL** A value boundary above which implementers are free to define
16 their own directive values.

17 3.2.21 IO Forwarding Channels

18 *PMIx v3.0* The `pmix_iof_channel_t` structure is a `uint16_t` type that defines a set of bit-mask flags
19 for specifying IO forwarding channels. These can be bitwise OR’d together to reference multiple
20 channels.

21 **PMIX_FWD_NO_CHANNELS** Forward no channels

22 **PMIX_FWD_STDIN_CHANNEL** Forward stdin

23 **PMIX_FWD_STDOUT_CHANNEL** Forward stdout

24 **PMIX_FWD_STDERR_CHANNEL** Forward stderr

25 **PMIX_FWD_STDDIAG_CHANNEL** Forward stddiag, if available

26 **PMIX_FWD_ALL_CHANNELS** Forward all available channels

27 3.2.22 Environmental Variable Structure

28 *PMIx v3.0* Define a structure for specifying environment variable modifications. Standard environment
29 variables (e.g., **PATH**, **LD_LIBRARY_PATH**, and **LD_PRELOAD**) take multiple arguments
30 separated by delimiters. Unfortunately, the delimiters depend upon the variable itself - some use
31 semi-colons, some colons, etc. Thus, the operation requires not only the name of the variable to be
32 modified and the value to be inserted, but also the separator to be used when composing the
33 aggregate value.

```

1 typedef struct
2     char *envar;
3     char *value;
4     char separator;
5     pmix_envar_t;

```

6 3.2.23 Environmental variable support macros

7 The following macros are provided to support the `pmix_envar_t` structure.

8 3.2.23.1 Initialize the `pmix_envar_t` structure

9 Initialize the `pmix_envar_t` fields

10 *PMIx v3.0*

```

10 PMIX_ENVAR_CONSTRUCT (m)
11 IN m
12     Pointer to the structure to be initialized (pointer to pmix_envar_t)

```

13 3.2.23.2 Destruct the `pmix_envar_t` structure

14 Clear the `pmix_envar_t` fields

15 *PMIx v3.0*

```

15 PMIX_ENVAR_DESTRUCT (m)
16 IN m
17     Pointer to the structure to be destructed (pointer to pmix_envar_t)

```

18 3.2.23.3 Create a `pmix_envar_t` array

19 Allocate and initialize an array of `pmix_envar_t` structures

20 *PMIx v3.0*

```

20 PMIX_ENVAR_CREATE (m, n)
21 INOUT m
22     Address where the pointer to the array of pmix_envar_t structures shall be stored (handle)
23 IN n
24     Number of structures to be allocated (size_t)

```

1 3.2.23.4 Free a `pmix_envar_t` array

2 Release an array of `pmix_envar_t` structures

PMIx v3.0

C

3 `PMIX_ENVAR_FREE(m, n)`

C

4 **IN** `m`

Pointer to the array of `pmix_envar_t` structures (handle)

6 **IN** `n`

Number of structures in the array (`size_t`)

8 3.2.23.5 Load a `pmix_envar_t` structure

9 Load values into a `pmix_envar_t`

PMIx v2.0

C

10 `PMIX_ENVAR_LOAD(m, e, v, s)`

C

11 **IN** `m`

Pointer to the structure to be loaded (pointer to `pmix_envar_t`)

13 **IN** `e`

Environmental variable name (`char*`)

15 **IN** `v`

Value of variable (`char*`)

17 **IN** `v`

Separator character (`char`)

19 3.2.24 Lookup Returned Data Structure

20 The `pmix_pdata_t` structure is used by `PMIx_Lookup` to describe the data being accessed.

PMIx v1.0

C

```
21 typedef struct pmix_pdata {  
22     pmix_proc_t proc;  
23     pmix_key_t key;  
24     pmix_value_t value;  
25 } pmix_pdata_t;
```

C

26 3.2.25 Lookup data structure support macros

27 The following macros are provided to support the `pmix_pdata_t` structure.

1 3.2.25.1 Initialize the `pmix_pdata_t` structure

2 Initialize the `pmix_pdata_t` fields

PMIx v1.0

▼  C 

3 **PMIX_PDATA_CONSTRUCT** (m)

▲  C 

4 **IN** m

5 Pointer to the structure to be initialized (pointer to `pmix_pdata_t`)

6 3.2.25.2 Destruct the `pmix_pdata_t` structure

7 Destruct the `pmix_pdata_t` fields

PMIx v1.0

▼  C 

8 **PMIX_PDATA_DESTRUCT** (m)

▲  C 

9 **IN** m

10 Pointer to the structure to be destructed (pointer to `pmix_pdata_t`)

11 3.2.25.3 Create a `pmix_pdata_t` array

12 Allocate and initialize an array of `pmix_pdata_t` structures

PMIx v1.0

▼  C 

13 **PMIX_PDATA_CREATE** (m, n)

▲  C 

14 **INOUT** m

15 Address where the pointer to the array of `pmix_pdata_t` structures shall be stored (handle)

16 **IN** n

17 Number of structures to be allocated (**size_t**)

18 3.2.25.4 Free a `pmix_pdata_t` array

19 Release an array of `pmix_pdata_t` structures

PMIx v1.0

▼  C 

20 **PMIX_PDATA_FREE** (m, n)

▲  C 

21 **IN** m

22 Pointer to the array of `pmix_pdata_t` structures (handle)

23 **IN** n

24 Number of structures in the array (**size_t**)

1 3.2.25.5 Load a lookup data structure

2 Summary

3 Load key, process identifier, and data value into a `pmix_pdata_t` structure.

PMIx v1.0

4 `PMIX_PDATA_LOAD(m, p, k, d, t);`

5 **IN** `m`

6 Pointer to the `pmix_pdata_t` structure into which the key and data are to be loaded
7 (pointer to `pmix_pdata_t`)

8 **IN** `p`

9 Pointer to the `pmix_proc_t` structure containing the identifier of the process being
10 referenced (pointer to `pmix_proc_t`)

11 **IN** `k`

12 String key to be loaded - must be less than or equal to `PMIX_MAX_KEYLEN` in length
13 (handle)

14 **IN** `d`

15 Pointer to the data value to be loaded (handle)

16 **IN** `t`

17 Type of the provided data value (`pmix_data_type_t`)

18 This macro simplifies the loading of key, process identifier, and data into a `pmix_proc_t` by
19 correctly assigning values to the structure's fields.

Advice to users

20 Key, process identifier, and data will all be copied into the `pmix_pdata_t` - thus, the source
21 information can be modified or free'd without affecting the copied data once the macro has
22 completed.

1 3.2.25.6 Transfer a lookup data structure

2 Summary

3 Transfer key, process identifier, and data value between two `pmix_pdata_t` structures.

PMIx v2.0

C

```
4 PMIX_PDATA_XFER(d, s);
```

C

5 **IN** d

6 Pointer to the destination `pmix_pdata_t` (pointer to `pmix_pdata_t`)

7 **IN** s

8 Pointer to the source `pmix_pdata_t` (pointer to `pmix_pdata_t`)

9 This macro simplifies the transfer of key and data between two `pmix_pdata_t` structures.

Advice to users

10 Key, process identifier, and data will all be copied into the destination `pmix_pdata_t` - thus, the
11 source `pmix_pdata_t` may free'd without affecting the copied data once the macro has
12 completed.

13 3.2.26 Application Structure

14 The `pmix_app_t` structure describes the application context for the `PMIx_Spawn` and
15 `PMIx_Spawn_nb` operations.

PMIx v1.0

C

```
16 typedef struct pmix_app {  
17     /** Executable */  
18     char *cmd;  
19     /** Argument set, NULL terminated */  
20     char **argv;  
21     /** Environment set, NULL terminated */  
22     char **env;  
23     /** Current working directory */  
24     char *cwd;  
25     /** Maximum processes with this profile */  
26     int maxprocs;  
27     /** Array of info keys describing this application*/  
28     pmix_info_t *info;  
29     /** Number of info keys in 'info' array */  
30     size_t ninfo;  
31 } pmix_app_t;
```

C

1 3.2.27 App structure support macros

2 The following macros are provided to support the `pmix_app_t` structure.

3 3.2.27.1 Initialize the `pmix_app_t` structure

4 Initialize the `pmix_app_t` fields

PMIx v1.0 ▼ `PMIX_APP_CONSTRUCT` (m) C

5 `PMIX_APP_CONSTRUCT` (m)

6 **IN** m

7 Pointer to the structure to be initialized (pointer to `pmix_app_t`)

8 3.2.27.2 Destruct the `pmix_app_t` structure

9 Destruct the `pmix_app_t` fields

PMIx v1.0 ▼ `PMIX_APP_DESTRUCT` (m) C

10 `PMIX_APP_DESTRUCT` (m)

11 **IN** m

12 Pointer to the structure to be destructed (pointer to `pmix_app_t`)

13 3.2.27.3 Create a `pmix_app_t` array

14 Allocate and initialize an array of `pmix_app_t` structures

PMIx v1.0 ▼ `PMIX_APP_CREATE` (m, n) C

15 `PMIX_APP_CREATE` (m, n)

16 **INOUT** m

17 Address where the pointer to the array of `pmix_app_t` structures shall be stored (handle)

18 **IN** n

19 Number of structures to be allocated (`size_t`)

20 3.2.27.4 Free a `pmix_app_t` array

21 Release an array of `pmix_app_t` structures

PMIx v1.0 ▼ `PMIX_APP_FREE` (m, n) C

22 `PMIX_APP_FREE` (m, n)

23 **IN** m

24 Pointer to the array of `pmix_app_t` structures (handle)

25 **IN** n

26 Number of structures in the array (`size_t`)

1 3.2.27.5 Create the `pmix_info_t` array of application directives

2 Create an array of `pmix_info_t` structures for passing application-level directives, updating the
3 `ninfo` field of the `pmix_app_t` structure.

PMIx v2.2

▼ C _____ ▼

4 **PMIX_APP_INFO_CREATE** (m, n)

▲ C _____ ▲

5 **IN** m
6 Pointer to the `pmix_app_t` structure (handle)

7 **IN** n
8 Number of directives to be allocated (**size_t**)

9 3.2.28 Query Structure

10 The `pmix_query_t` structure is used by `PMIx_Query_info_nb` to describe a single query
11 operation.

PMIx v2.0

▼ C _____ ▼

```
12 typedef struct pmix_query {  
13     char **keys;  
14     pmix_info_t *qualifiers;  
15     size_t nqual;  
16 } pmix_query_t;
```

▲ C _____ ▲

17 3.2.29 Query structure support macros

18 The following macros are provided to support the `pmix_query_t` structure.

19 3.2.29.1 Initialize the `pmix_query_t` structure

20 Initialize the `pmix_query_t` fields

PMIx v2.0

▼ C _____ ▼

21 **PMIX_QUERY_CONSTRUCT** (m)

▲ C _____ ▲

22 **IN** m
23 Pointer to the structure to be initialized (pointer to `pmix_query_t`)

1 3.2.29.2 Destruct the `pmix_query_t` structure

2 Destruct the `pmix_query_t` fields

PMIx v2.0

▼ `PMIX_QUERY_DESTRUCT` C

3 `PMIX_QUERY_DESTRUCT` (m)

▲ C

4 **IN** m

5 Pointer to the structure to be destructed (pointer to `pmix_query_t`)

6 3.2.29.3 Create a `pmix_query_t` array

7 Allocate and initialize an array of `pmix_query_t` structures

PMIx v2.0

▼ `PMIX_QUERY_CREATE` C

8 `PMIX_QUERY_CREATE` (m, n)

▲ C

9 **INOUT** m

10 Address where the pointer to the array of `pmix_query_t` structures shall be stored (handle)

11 **IN** n

12 Number of structures to be allocated (`size_t`)

13 3.2.29.4 Free a `pmix_query_t` array

14 Release an array of `pmix_query_t` structures

PMIx v2.0

▼ `PMIX_QUERY_FREE` C

15 `PMIX_QUERY_FREE` (m, n)

▲ C

16 **IN** m

17 Pointer to the array of `pmix_query_t` structures (handle)

18 **IN** n

19 Number of structures in the array (`size_t`)

20 3.2.29.5 Create the `pmix_info_t` array of query qualifiers

21 Create an array of `pmix_info_t` structures for passing query qualifiers, updating the `nqual` field
22 of the `pmix_query_t` structure.

PMIx v2.2

▼ `PMIX_QUERY_QUALIFIERS_CREATE` C

23 `PMIX_QUERY_QUALIFIERS_CREATE` (m, n)

▲ C

24 **IN** m

25 Pointer to the `pmix_query_t` structure (handle)

26 **IN** n

27 Number of qualifiers to be allocated (`size_t`)

1 3.2.30 Attribute registration structure

2 The `pmix_regattr_t` structure is used to register attribute support for a PMIx function.

PMIx v4.0

```
3 typedef struct pmix_regattr {  
4     char *name;  
5     pmix_key_t *string;  
6     pmix_data_type_t type;  
7     pmix_info_t *info;  
8     size_t ninfo;  
9     char **description;  
10 } pmix_regattr_t;
```

11 Note that in this structure:

- 12 • the *name* is the actual name of the attribute - e.g., "PMIX_MAX_PROCS"; and
- 13 • the *string* is the literal string value of the attribute - e.g., "pmix.max.size" for the
14 `PMIX_MAX_PROCS` attribute
- 15 • *type* must be a PMIx data type identifying the type of data associated with this attribute.
- 16 • the *info* array contains machine-usable information regarding the range of accepted values. This
17 may include entries for `PMIX_MIN_VALUE`, `PMIX_MAX_VALUE`, `PMIX_ENUM_VALUE`, or
18 a combination of them. For example, an attribute that supports all positive integers might
19 delineate it by including a `pmix_info_t` with a key of `PMIX_MIN_VALUE`, type of
20 `PMIX_INT`, and value of zero. The lack of an entry for `PMIX_MAX_VALUE` indicates that
21 there is no ceiling to the range of accepted values.
- 22 • *ninfo* indicates the number of elements in the *info* array
- 23 • The *description* field consists of a **NULL**-terminated array of strings describing the attribute,
24 optionally including a human-readable description of the range of accepted values - e.g., "ALL
25 POSITIVE INTEGERS", or a comma-delimited list of enum value names. No correlation
26 between the number of entries in the *description* and the number of elements in the *info* array is
27 implied or required.

28 The attribute *name* and *string* fields must be **NULL**-terminated strings composed of standard
29 alphanumeric values supported by common utilities such as *strcmp*.

Advice to PMIx library implementers

30 Although not strictly required, PMIx library implementers are strongly encouraged to provide both
31 human-readable and machine-parsable descriptions of supported attributes.

1 Although not strictly required, host environments are strongly encouraged to provide both
 2 human-readable and machine-parsable descriptions of supported attributes when registering them.

3 3.2.31 Attribute registration structure support macros

4 The following macros are provided to support the `pmix_regattr_t` structure.

5 3.2.31.1 Initialize the `pmix_regattr_t` structure

6 Initialize the `pmix_regattr_t` fields

PMIx v4.0

7 **PMIX_REGATTR_CONSTRUCT** (m)

8 **IN** m

9 Pointer to the structure to be initialized (pointer to `pmix_regattr_t`)

10 3.2.31.2 Destruct the `pmix_regattr_t` structure

11 Destruct the `pmix_regattr_t` fields, releasing all strings.

PMIx v4.0

12 **PMIX_REGATTR_DESTRUCT** (m)

13 **IN** m

14 Pointer to the structure to be destructed (pointer to `pmix_regattr_t`)

15 3.2.31.3 Create a `pmix_regattr_t` array

16 Allocate and initialize an array of `pmix_regattr_t` structures

PMIx v4.0

17 **PMIX_REGATTR_CREATE** (m, n)

18 **INOUT** m

19 Address where the pointer to the array of `pmix_regattr_t` structures shall be stored
 20 (handle)

21 **IN** n

22 Number of structures to be allocated (`size_t`)

1 3.2.31.4 Free a `pmix_regattr_t` array

2 Release an array of `pmix_regattr_t` structures

PMIx v4.0

C

3 **PMIX_REGATTR_FREE**(*m*, *n*)

C

4 **INOUT** *m*

5 Pointer to the array of `pmix_regattr_t` structures (handle)

6 **IN** *n*

7 Number of structures in the array (**size_t**)

8 3.2.31.5 Load a `pmix_regattr_t` structure

9 Load values into a `pmix_regattr_t` structure. The macro can be called multiple times to add
10 as many strings as desired to the same structure by passing the same address and a **NULL** key to the
11 macro. Note that the *t* type value must be given each time.

PMIx v4.0

C

12 **PMIX_REGATTR_LOAD**(*a*, *n*, *k*, *t*, *ni*, *v*)

C

13 **IN** *a*

14 Pointer to the structure to be loaded (pointer to `pmix_proc_t`)

15 **IN** *n*

16 String name of the attribute (string)

17 **IN** *k*

18 Key value to be loaded (`pmix_key_t`)

19 **IN** *t*

20 Type of data associated with the provided key (`pmix_data_type_t`)

21 **IN** *ni*

22 Number of `pmix_info_t` elements to be allocated in *info* (**size_t**)

23 **IN** *v*

24 One-line description to be loaded (more can be added separately) (string)

25 3.2.31.6 Transfer a `pmix_regattr_t` to another `pmix_regattr_t`

26

27 Non-destructively transfer the contents of a `pmix_regattr_t` structure to another one.

PMIx v4.0

C

28 **PMIX_REGATTR_XFER**(*m*, *n*)

C

29 **INOUT** *m*

30 Pointer to the destination `pmix_regattr_t` structure (handle)

31 **IN** *n*

32 Pointer to the source `pmix_regattr_t` structure (handle)

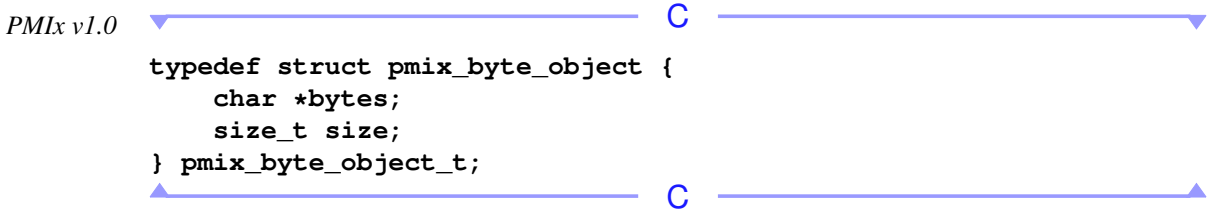
1 3.2.32 PMIx Group Directives

2 *PMIx v4.0* The `pmix_group_opt_t` type is an enumerated type used with the `PMIx_Group_join` API
3 to indicate *accept* or *decline* of the invitation - these are provided for readability of user code:

- 4 `PMIX_GROUP_DECLINE` Decline the invitation
- 5 `PMIX_GROUP_ACCEPT` Accept the invitation.

6 3.2.33 Byte Object Type

7 The `pmix_byte_object_t` structure describes a raw byte sequence.

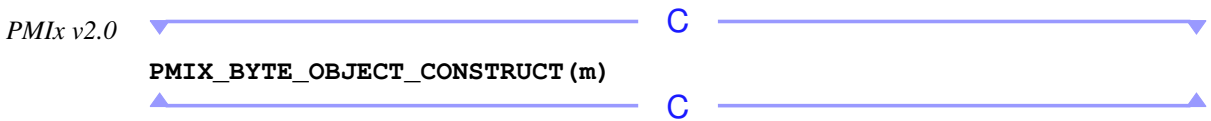
```
PMIx v1.0 
8 typedef struct pmix_byte_object {
9     char *bytes;
10    size_t size;
11 } pmix_byte_object_t;
```

12 3.2.34 Byte object support macros

13 The following macros support the `pmix_byte_object_t` structure.

14 3.2.34.1 Initialize the `pmix_byte_object_t` structure

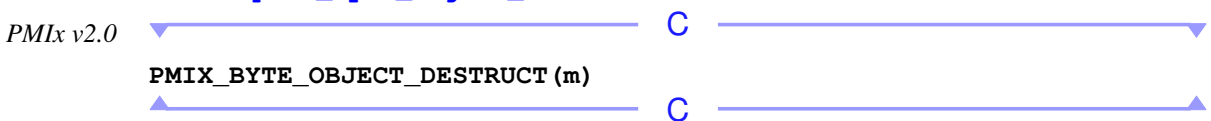
15 Initialize the `pmix_byte_object_t` fields

```
PMIx v2.0 
16 PMIX_BYTE_OBJECT_CONSTRUCT (m)
```

17 **IN** m
18 Pointer to the structure to be initialized (pointer to `pmix_byte_object_t`)

19 3.2.34.2 Destruct the `pmix_byte_object_t` structure

20 Clear the `pmix_byte_object_t` fields

```
PMIx v2.0 
21 PMIX_BYTE_OBJECT_DESTRUCT (m)
```

22 **IN** m
23 Pointer to the structure to be destructed (pointer to `pmix_byte_object_t`)

1 3.2.34.3 Create a `pmix_byte_object_t` structure

2 Allocate and initialize an array of `pmix_byte_object_t` structures

PMIx v2.0

▼ `C` _____ ▼

3 **PMIX_BYTE_OBJECT_CREATE** (*m*, *n*)

▲ _____ `C` _____ ▲

4 **INOUT** *m*

5 Address where the pointer to the array of `pmix_byte_object_t` structures shall be
6 stored (handle)

7 **IN** *n*

8 Number of structures to be allocated (**size_t**)

9 3.2.34.4 Free a `pmix_byte_object_t` array

10 Release an array of `pmix_byte_object_t` structures

PMIx v2.0

▼ `C` _____ ▼

11 **PMIX_BYTE_OBJECT_FREE** (*m*, *n*)

▲ _____ `C` _____ ▲

12 **IN** *m*

13 Pointer to the array of `pmix_byte_object_t` structures (handle)

14 **IN** *n*

15 Number of structures in the array (**size_t**)

16 3.2.34.5 Load a `pmix_byte_object_t` structure

17 Load values into a `pmix_byte_object_t`

PMIx v2.0

▼ `C` _____ ▼

18 **PMIX_BYTE_OBJECT_LOAD** (*b*, *d*, *s*)

▲ _____ `C` _____ ▲

19 **IN** *b*

20 Pointer to the structure to be loaded (pointer to `pmix_byte_object_t`)

21 **IN** *d*

22 Pointer to the data to be loaded (**char***)

23 **IN** *s*

24 Number of bytes in the data array (**size_t**)

1 3.2.35 Data Array Structure

2 The `pmix_data_array_t` structure defines an array data structure.

PMIx v2.0

```
3 typedef struct pmix_data_array {  
4     pmix_data_type_t type;  
5     size_t size;  
6     void *array;  
7 } pmix_data_array_t;
```

8 3.2.36 Data array support macros

9 The following macros support the `pmix_data_array_t` structure.

10 3.2.36.1 Initialize a `pmix_data_array_t` structure

11 Initialize the `pmix_data_array_t` fields, allocating memory for the array of the indicated type.

PMIx v2.2

```
12 PMIX_DATA_ARRAY_CONSTRUCT(m, n, t)
```

13 **IN** m

14 Pointer to the structure to be initialized (pointer to `pmix_data_array_t`)

15 **IN** n

16 Number of elements in the array (`size_t`)

17 **IN** t

18 PMIx data type of the array elements (`pmix_data_type_t`)

19 3.2.36.2 Destruct a `pmix_data_array_t` structure

20 Destruct the `pmix_data_array_t`, releasing the memory in the array.

PMIx v2.2

```
21 PMIX_DATA_ARRAY_DESTRUCT(m)
```

22 **IN** m

23 Pointer to the structure to be destructed (pointer to `pmix_data_array_t`)

1 3.2.36.3 Create a `pmix_data_array_t` structure

2 Allocate memory for the `pmix_data_array_t` object itself, and then allocate memory for the
3 array of the indicated type.

PMIx v2.2

```
▼ _____ C _____ ▼  
4 PMIX_DATA_ARRAY_CREATE(m, n, t)  
▲ _____ C _____ ▲
```

5 **INOUT** m

6 Variable to be set to the address of the structure (pointer to `pmix_data_array_t`)

7 **IN** n

8 Number of elements in the array (`size_t`)

9 **IN** t

10 PMIx data type of the array elements (`pmix_data_type_t`)

11 3.2.36.4 Free a `pmix_data_array_t` structure

12 Release the memory in the array, and then release the `pmix_data_array_t` object itself.

PMIx v2.2

```
▼ _____ C _____ ▼  
13 PMIX_DATA_ARRAY_FREE(m)  
▲ _____ C _____ ▲
```

14 **IN** m

15 Pointer to the structure to be released (pointer to `pmix_data_array_t`)

16 3.2.37 Argument Array Macros

17 The following macros support the construction and release of **NULL**-terminated argv arrays of
18 strings.

19 3.2.37.1 Argument array extension

20 **Summary**

21 Append a string to a **NULL**-terminated, argv-style array of strings.

```
▼ _____ C _____ ▼  
22 PMIX_ARGV_APPEND(r, a, b);  
▲ _____ C _____ ▲
```

23 **OUT** r

24 Status code indicating success or failure of the operation (`pmix_status_t`)

25 **INOUT** a

26 Argument list (pointer to **NULL**-terminated array of strings)

27 **IN** b

28 Argument to append to the list (string)

Description

This function helps the caller build the `argv` portion of `pmix_app_t` structure, arrays of keys for querying, or other places where argv-style string arrays are required in the way that the PRI expects it to be constructed.

Advice to users

The provided argument is copied into the destination array - thus, the source string can be free'd without affecting the array once the macro has completed.

3.2.37.2 Argument array extension - unique

Summary

Append a string to a NULL-terminated, argv-style array of strings, but only if the provided argument doesn't already exist somewhere in the array.

```
PMIX_ARGV_APPEND_UNIQUE(r, a, b);
```

OUT `r`

Status code indicating success or failure of the operation (`pmix_status_t`)

INOUT `a`

Argument list (pointer to NULL-terminated array of strings)

IN `b`

Argument to append to the list (string)

Description

This function helps the caller build the `argv` portion of `pmix_app_t` structure, arrays of keys for querying, or other places where argv-style string arrays are required in the way that the PRI expects it to be constructed.

Advice to users

The provided argument is copied into the destination array - thus, the source string can be free'd without affecting the array once the macro has completed.

1 3.2.37.3 Argument array release

2 Summary

3 Free an argv-style array and all of the strings that it contains

▼ C ▼

4 **PMIX_ARGV_FREE (a) ;**

▲ C ▲

5 **IN a**

6 Argument list (pointer to NULL-terminated array of strings)

7 Description

8 This function releases the array and all of the strings it contains.

9 3.2.37.4 Argument array split

10 Summary

11 Split a string into a NULL-terminated argv array.

▼ C ▼

12 **PMIX_ARGV_SPLIT (a, b, c) ;**

▲ C ▲

13 **OUT a**

14 Resulting argv-style array (**char****)

15 **IN b**

16 String to be split (**char***)

17 **IN c**

18 Delimiter character (**char**)

19 Description

20 Split an input string into a NULL-terminated argv array. Do not include empty strings in the
21 resulting array.

▼ Advice to users ▼

22 All strings are inserted into the argv array by value; the newly-allocated array makes no references
23 to the src_string argument (i.e., it can be freed after calling this function without invalidating the
24 output argv array)

▲

1 3.2.37.5 Argument array join

2 Summary

3 Join all the elements of an argv array into a single newly-allocated string.

▼  ▼

4 `PMIX_ARGV_JOIN(a, b, c);`

▲  ▲

5 **OUT** a

6 Resulting string (**char***)

7 **IN** b

8 Argv-style array to be joined (**char****)

9 **IN** c

10 Delimiter character (**char**)

11 Description

12 Join all the elements of an argv array into a single newly-allocated string.

13 3.2.37.6 Argument array count

14 Summary

15 Return the length of a NULL-terminated argv array.

▼  ▼

16 `PMIX_ARGV_COUNT(r, a);`

▲  ▲

17 **OUT** r

18 Number of strings in the array (integer)

19 **IN** a

20 Argv-style array (**char****)

21 Description

22 Count the number of elements in an argv array

23 3.2.37.7 Argument array copy

24 Summary

25 Copy an argv array, including copying all off its strings.

▼  ▼

26 `PMIX_ARGV_COPY(a, b);`

▲  ▲

27 **OUT** a

28 New argv-style array (**char****)

29 **IN** b

30 Argv-style array (**char****)

1 **Description**
2 Copy an argv array, including copying all off its strings.

3 **3.2.38 Set Environment Variable**

4 **Summary**
5 Set an environment variable in a **NULL**-terminated, env-style array

▼── C ───▶
6 **PMIX_SETENV**(*r*, *name*, *value*, *env*);
▲── C ───▶

7 **OUT** *r*
8 Status code indicating success or failure of the operation ([pmix_status_t](#))
9 **IN** *name*
10 Argument name (string)
11 **IN** *value*
12 Argument value (string)
13 **INOUT** *env*
14 Environment array to update (pointer to array of strings)

15 **Description**
16 Similar to **setenv** from the C API, this allows the caller to set an environment variable in the
17 specified **env** array, which could then be passed to the [pmix_app_t](#) structure or any other
18 destination.

▼── Advice to users ───▶
19 The provided name and value are copied into the destination environment array - thus, the source
20 strings can be free'd without affecting the array once the macro has completed.
▲──▶

1 3.3 Generalized Data Types Used for Packing/Unpacking

2 The `pmix_data_type_t` structure is a `uint16_t` type for identifying the data type for
3 packing/unpacking purposes. New data type values introduced in this version of the Standard are
4 shown in **magenta**.

Advice to PMIx library implementers

5 The following constants can be used to set a variable of the type `pmix_data_type_t`. Data
6 types in the PMIx Standard are defined in terms of the C-programming language. Implementers
7 wishing to support other languages should provide the equivalent definitions in a
8 language-appropriate manner. Additionally, a PMIx implementation may choose to add additional
9 types.

10	<code>PMIX_UNDEF</code>	Undefined
11	<code>PMIX_BOOL</code>	Boolean (converted to/from native <code>true/false</code>) (<code>bool</code>)
12	<code>PMIX_BYTE</code>	A byte of data (<code>uint8_t</code>)
13	<code>PMIX_STRING</code>	<code>NULL</code> terminated string (<code>char*</code>)
14	<code>PMIX_SIZE</code>	Size <code>size_t</code>
15	<code>PMIX_PID</code>	Operating process identifier (PID) (<code>pid_t</code>)
16	<code>PMIX_INT</code>	Integer (<code>int</code>)
17	<code>PMIX_INT8</code>	8-byte integer (<code>int8_t</code>)
18	<code>PMIX_INT16</code>	16-byte integer (<code>int16_t</code>)
19	<code>PMIX_INT32</code>	32-byte integer (<code>int32_t</code>)
20	<code>PMIX_INT64</code>	64-byte integer (<code>int64_t</code>)
21	<code>PMIX_UINT</code>	Unsigned integer (<code>unsigned int</code>)
22	<code>PMIX_UINT8</code>	Unsigned 8-byte integer (<code>uint8_t</code>)
23	<code>PMIX_UINT16</code>	Unsigned 16-byte integer (<code>uint16_t</code>)
24	<code>PMIX_UINT32</code>	Unsigned 32-byte integer (<code>uint32_t</code>)
25	<code>PMIX_UINT64</code>	Unsigned 64-byte integer (<code>uint64_t</code>)
26	<code>PMIX_FLOAT</code>	Float (<code>float</code>)
27	<code>PMIX_DOUBLE</code>	Double (<code>double</code>)
28	<code>PMIX_TIMEVAL</code>	Time value (<code>struct timeval</code>)
29	<code>PMIX_TIME</code>	Time (<code>time_t</code>)
30	<code>PMIX_STATUS</code>	Status code <code>pmix_status_t</code>
31	<code>PMIX_VALUE</code>	Value (<code>pmix_value_t</code>)
32	<code>PMIX_PROC</code>	Process (<code>pmix_proc_t</code>)
33	<code>PMIX_APP</code>	Application context
34	<code>PMIX_INFO</code>	Info object
35	<code>PMIX_PDATA</code>	Pointer to data
36	<code>PMIX_BUFFER</code>	Buffer
37	<code>PMIX_BYTE_OBJECT</code>	Byte object (<code>pmix_byte_object_t</code>)
38	<code>PMIX_KVAL</code>	Key/value pair

```

1   PMIX_PERSIST    Persistence ( pmix\_persistence\_t )
2   PMIX_POINTER   Pointer to an object (void*)
3   PMIX_SCOPE     Scope ( pmix\_scope\_t )
4   PMIX_DATA_RANGE Range for data ( pmix\_data\_range\_t )
5   PMIX_COMMAND   PMIx command code (used internally)
6   PMIX_INFO_DIRECTIVES Directives flag for pmix\_info\_t (
7       pmix\_info\_directives\_t )
8   PMIX_DATA_TYPE Data type code ( pmix\_data\_type\_t )
9   PMIX_PROC_STATE Process state ( pmix\_proc\_state\_t )
10  PMIX_PROC_INFO Process information ( pmix\_proc\_info\_t )
11  PMIX_DATA_ARRAY Data array ( pmix\_data\_array\_t )
12  PMIX_PROC_RANK Process rank ( pmix\_rank\_t )
13  PMIX_QUERY     Query structure ( pmix\_query\_t )
14  PMIX_COMPRESSED_STRING String compressed with zlib (char*)
15  PMIX_ALLOC_DIRECTIVE Allocation directive ( pmix\_alloc\_directive\_t )
16  PMIX_IOF_CHANNEL Input/output forwarding channel ( pmix\_iof\_channel\_t )
17  PMIX_ENVAR     Environmental variable structure ( pmix\_envar\_t )
18  PMIX_COORD    Structure containing network coordinates ( pmix\_coord\_t )
19  PMIX_REGATTR  Structure supporting attribute registrations ( pmix\_regattr\_t )
20  PMIX_REGEX    Regular expressions - can be a valid NULL-terminated string or an arbitrary
21      array of bytes

```

22 3.4 Reserved attributes

23 The PMIx standard defines a relatively small set of APIs and the caller may customize the behavior
24 of the API by passing one or more attributes to that API. Additionally, attributes may be keys
25 passed to [PMIx_Get](#) calls to access the specified values from the system.

26 Each attribute is represented by a *key* string, and a type for the associated *value*. This section
27 defines a set of **reserved** keys which are prefixed with **pmix.** to designate them as PMIx standard
28 reserved keys. All definitions were introduced in version 1 of the standard unless otherwise marked.

29 Applications or associated libraries (e.g., MPI) may choose to define additional attributes. The
30 attributes defined in this section are of the system and job as opposed to the attributes that the
31 application (or associated libraries) might choose to expose. Due to this extensibility the
32 [PMIx_Get](#) API will return [PMIX_ERR_NOT_FOUND](#) if the provided *key* cannot be found.

33 Attributes added in this version of the standard are shown in *magenta* to distinguish them from
34 those defined in prior versions, which are shown in *black*. Deprecated attributes are shown in *green*
35 and will be removed in future versions of the standard.

36 **PMIX_ATTR_UNDEF NULL (NULL)**
37 Constant representing an undefined attribute.

1 3.4.1 Initialization attributes

2 These attributes are defined to assist the caller with initialization by passing them into the
3 appropriate initialization API - thus, they are not typically accessed via the `PMIx_Get` API.

4 `PMIX_EVENT_BASE` "`pmix.evbase`" (`struct event_base *`)
5 Pointer to libevent¹ `event_base` to use in place of the internal progress thread.
6 `PMIX_SERVER_TOOL_SUPPORT` "`pmix.srvr.tool`" (`bool`)
7 The host RM wants to declare itself as willing to accept tool connection requests.
8 `PMIX_SERVER_REMOTE_CONNECTIONS` "`pmix.srvr.remote`" (`bool`)
9 Allow connections from remote tools. Forces the PMIx server to not exclusively use
10 loopback device.
11 `PMIX_SERVER_SYSTEM_SUPPORT` "`pmix.srvr.sys`" (`bool`)
12 The host RM wants to declare itself as being the local system server for PMIx connection
13 requests.
14 `PMIX_SERVER_TMPDIR` "`pmix.srvr.tmpdir`" (`char*`)
15 Top-level temporary directory for all client processes connected to this server, and where the
16 PMIx server will place its tool rendezvous point and contact information.
17 `PMIX_SYSTEM_TMPDIR` "`pmix.sys.tmpdir`" (`char*`)
18 Temporary directory for this system, and where a PMIx server that declares itself to be a
19 system-level server will place a tool rendezvous point and contact information.
20 `PMIX_SERVER_ENABLE_MONITORING` "`pmix.srv.monitor`" (`bool`)
21 Enable PMIx internal monitoring by the PMIx server.
22 `PMIX_SERVER_NAMESPACE` "`pmix.srv.namespace`" (`char*`)
23 Name of the namespace to use for this PMIx server.
24 `PMIX_SERVER_RANK` "`pmix.srv.rank`" (`pmix_rank_t`)
25 Rank of this PMIx server
26 `PMIX_SERVER_GATEWAY` "`pmix.srv.gway`" (`bool`)
27 Server is acting as a gateway for PMIx requests that cannot be serviced on backend nodes
28 (e.g., logging to email)

29 3.4.2 Tool-related attributes

30 These attributes are defined to assist PMIx-enabled tools to connect with the PMIx server by
31 passing them into the `PMIx_tool_init` API - thus, they are not typically accessed via the
32 `PMIx_Get` API.

33 `PMIX_TOOL_NAMESPACE` "`pmix.tool.namespace`" (`char*`)
34 Name of the namespace to use for this tool.
35 `PMIX_TOOL_RANK` "`pmix.tool.rank`" (`uint32_t`)
36 Rank of this tool.
37 `PMIX_SERVER_PIDINFO` "`pmix.srvr.pidinfo`" (`pid_t`)
38 PID of the target PMIx server for a tool.
39 `PMIX_CONNECT_TO_SYSTEM` "`pmix.cnct.sys`" (`bool`)

¹<http://libevent.org/>

1 The requestor requires that a connection be made only to a local, system-level PMIx server.

2 **PMIX_CONNECT_SYSTEM_FIRST** "pmix.cnct.sys.first" (bool)

3 Preferentially, look for a system-level PMIx server first.

4 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)

5 uniform resource identifier (URI) of the PMIx server to be contacted.

6 **PMIX_SERVER_HOSTNAME** "pmix.srvr.host" (char*)

7 Host where target PMIx server is located.

8 **PMIX_CONNECT_MAX_RETRIES** "pmix.tool.mretries" (uint32_t)

9 Maximum number of times to try to connect to PMIx server.

10 **PMIX_CONNECT_RETRY_DELAY** "pmix.tool.retry" (uint32_t)

11 Time in seconds between connection attempts to a PMIx server.

12 **PMIX_TOOL_DO_NOT_CONNECT** "pmix.tool.nocon" (bool)

13 The tool wants to use internal PMIx support, but does not want to connect to a PMIx server.

14 **PMIX_RECONNECT_SERVER** "pmix.tool.recon" (bool)

15 Tool is requesting to change server connections

16 **PMIX_LAUNCHER** "pmix.tool.launcher" (bool)

17 Tool is a launcher and needs rendezvous files created

18 3.4.3 Identification attributes

19 These attributes are defined to identify a process and it's associated PMIx-enabled library. They are
20 not typically accessed via the **PMIx_Get** API, and thus are not associated with a particular rank.

21 **PMIX_USERID** "pmix.euid" (uint32_t)

22 Effective user id.

23 **PMIX_GRPID** "pmix.egid" (uint32_t)

24 Effective group id.

25 **PMIX_DSTPATH** "pmix.dstpath" (char*)

26 Path to shared memory data storage (dstore) files.

27 **PMIX_VERSION_INFO** "pmix.version" (char*)

28 PMIx version of contractor.

29 **PMIX_REQUESTOR_IS_TOOL** "pmix.req.tool" (bool)

30 The requesting process is a PMIx tool.

31 **PMIX_REQUESTOR_IS_CLIENT** "pmix.req.client" (bool)

32 The requesting process is a PMIx client.

33 **PMIX_PSET_NAME** "pmix.pset.nm" (char*)

34 User-assigned name for the process set containing the given process.

1 3.4.4 Programming model attributes

2 These attributes are associated with programming models.

3 **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)
4 Programming model being initialized (e.g., "MPI" or "OpenMP")
5 **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)
6 Programming model implementation ID (e.g., "OpenMPI" or "MPICH")
7 **PMIX_MODEL_LIBRARY_VERSION** "pmix.mdl.vrs" (char*)
8 Programming model version string (e.g., "2.1.1")
9 **PMIX_THREADING_MODEL** "pmix.threads" (char*)
10 Threading model used (e.g., "pthreads")
11 **PMIX_MODEL_NUM_THREADS** "pmix.mdl.nthrds" (uint64_t)
12 Number of active threads being used by the model
13 **PMIX_MODEL_NUM_CPUS** "pmix.mdl.ncpu" (uint64_t)
14 Number of cpus being used by the model
15 **PMIX_MODEL_CPU_TYPE** "pmix.mdl.cputype" (char*)
16 Granularity - "hwthread", "core", etc.
17 **PMIX_MODEL_PHASE_NAME** "pmix.mdl.phase" (char*)
18 User-assigned name for a phase in the application execution (e.g., "cfd reduction")
19 **PMIX_MODEL_PHASE_TYPE** "pmix.mdl.ptype" (char*)
20 Type of phase being executed (e.g., "matrix multiply")
21 **PMIX_MODEL_AFFINITY_POLICY** "pmix.mdl.tap" (char*)
22 Thread affinity policy - e.g.: "master" (thread co-located with master thread), "close" (thread
23 located on cpu close to master thread), "spread" (threads load-balanced across available cpus)

24 3.4.5 UNIX socket rendezvous socket attributes

25 These attributes are used to describe a UNIX socket for rendezvous with the local RM by passing
26 them into the relevant initialization API - thus, they are not typically accessed via the [PMIx_Get](#)
27 API.

28 **PMIX_USOCK_DISABLE** "pmix.usock.disable" (bool)
29 Disable legacy UNIX socket (usock) support
30 **PMIX_SOCKET_MODE** "pmix.sockmode" (uint32_t)
31 POSIX *mode_t* (9 bits valid)
32 **PMIX_SINGLE_LISTENER** "pmix.sing.listnr" (bool)
33 Use only one rendezvous socket, letting priorities and/or environment parameters select the
34 active transport.

1 3.4.6 TCP connection attributes

2 These attributes are used to describe a TCP socket for rendezvous with the local RM by passing
3 them into the relevant initialization API - thus, they are not typically accessed via the [PMIx_Get](#)
4 API.

5 **PMIX_TCP_REPORT_URI** "pmix.tcp.repuri" (char*)

6 If provided, directs that the TCP URI be reported and indicates the desired method of
7 reporting: '-' for stdout, '+' for stderr, or filename.

8 **PMIX_TCP_URI** "pmix.tcp.uri" (char*)

9 The URI of the PMIx server to connect to, or a file name containing it in the form of
10 **file:<name of file containing it>**.

11 **PMIX_TCP_IF_INCLUDE** "pmix.tcp.ifinclude" (char*)

12 Comma-delimited list of devices and/or Classless Inter-Domain Routing (CIDR) notation to
13 include when establishing the TCP connection.

14 **PMIX_TCP_IF_EXCLUDE** "pmix.tcp.ifexclude" (char*)

15 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
16 TCP connection.

17 **PMIX_TCP_IPV4_PORT** "pmix.tcp.ipv4" (int)

18 The IPv4 port to be used.

19 **PMIX_TCP_IPV6_PORT** "pmix.tcp.ipv6" (int)

20 The IPv6 port to be used.

21 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)

22 Set to **true** to disable IPv4 family of addresses.

23 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)

24 Set to **true** to disable IPv6 family of addresses.

25 3.4.7 Global Data Storage (GDS) attributes

26 These attributes are used to define the behavior of the GDS used to manage key/value pairs by
27 passing them into the relevant initialization API - thus, they are not typically accessed via the
28 [PMIx_Get](#) API.

29 **PMIX_GDS_MODULE** "pmix.gds.mod" (char*)

30 Comma-delimited string of desired modules.

31 3.4.8 General process-level attributes

32 These attributes are used to define process attributes and are referenced by their process rank.

33 **PMIX_CPUSSET** "pmix.cpuset" (char*)

34 hwloc² bitmap to be applied to the process upon launch.

35 **PMIX_CREDENTIAL** "pmix.cred" (char*)

36 Security credential assigned to the process.

37 **PMIX_SPAWNED** "pmix.spawned" (bool)

²<https://www.open-mpi.org/projects/hwloc/>

1 true if this process resulted from a call to `PMIx_Spawn`.
2 `PMIX_ARCH` "pmix.arch" (uint32_t)
3 Architecture flag.

4 3.4.9 Scratch directory attributes

5 These attributes are used to define an application scratch directory and are referenced using the
6 `PMIX_RANK_WILDCARD` rank.

7 `PMIX_TMPDIR` "pmix.tmpdir" (char*)
8 Full path to the top-level temporary directory assigned to the session.
9 `PMIX_NSDIR` "pmix.nsdire" (char*)
10 Full path to the temporary directory assigned to the namespace, under `PMIX_TMPDIR`.
11 `PMIX_PROCDIR` "pmix.pdire" (char*)
12 Full path to the subdirectory under `PMIX_NSDIR` assigned to the process.
13 `PMIX_TDIR_RMCLEAN` "pmix.tdire.rmclean" (bool)
14 Resource Manager will clean session directories

15 3.4.10 Relative Rank Descriptive Attributes

16 These attributes are used to describe information about relative ranks as assigned by the RM, and
17 thus are referenced using the process rank except where noted.

18 `PMIX_CLUSTER_ID` "pmix.clid" (char*)
19 A string name for the cluster this proc is executing on
20 `PMIX_PROCID` "pmix.procid" (pmix_proc_t)
21 Process identifier
22 `PMIX_NAMESPACE` "pmix.namespace" (char*)
23 Namespace of the job.
24 `PMIX_JOBID` "pmix.jobid" (char*)
25 Job identifier assigned by the scheduler.
26 `PMIX_APPNUM` "pmix.appnum" (uint32_t)
27 Application number within the job.
28 `PMIX_RANK` "pmix.rank" (pmix_rank_t)
29 Process rank within the job.
30 `PMIX_GLOBAL_RANK` "pmix.grank" (pmix_rank_t)
31 Process rank spanning across all jobs in this session.
32 `PMIX_APP_RANK` "pmix.apprank" (pmix_rank_t)
33 Process rank within this application.
34 `PMIX_NPROC_OFFSET` "pmix.offset" (pmix_rank_t)
35 Starting global rank of this job - referenced using `PMIX_RANK_WILDCARD`.
36 `PMIX_LOCAL_RANK` "pmix.lrank" (uint16_t)
37 Local rank on this node within this job.
38 `PMIX_NODE_RANK` "pmix.nrank" (uint16_t)
39 Process rank on this node spanning all jobs.

1 **PMIX_LOCALLDR** "pmix.lldr" (pmix_rank_t)
2 Lowest rank on this node within this job - referenced using [PMIX_RANK_WILDCARD](#) .
3 **PMIX_APPLDR** "pmix.aldr" (pmix_rank_t)
4 Lowest rank in this application within this job - referenced using [PMIX_RANK_WILDCARD](#) .
5 **PMIX_PROC_PID** "pmix.ppid" (pid_t)
6 PID of specified process.
7 **PMIX_SESSION_ID** "pmix.session.id" (uint32_t)
8 Session identifier - referenced using [PMIX_RANK_WILDCARD](#) .
9 **PMIX_NODE_LIST** "pmix.nlist" (char*)
10 Comma-delimited list of nodes running processes for the specified namespace - referenced
11 using [PMIX_RANK_WILDCARD](#) .
12 **PMIX_ALLOCATED_NODELIST** "pmix.alist" (char*)
13 Comma-delimited list of all nodes in this allocation regardless of whether or not they
14 currently host processes - referenced using [PMIX_RANK_WILDCARD](#) .
15 **PMIX_HOSTNAME** "pmix.hname" (char*)
16 Name of the host where the specified process is running.
17 **PMIX_NODEID** "pmix.nodeid" (uint32_t)
18 Node identifier where the specified process is located, expressed as the node's index
19 (beginning at zero) in the array resulting from expansion of the [PMIX_NODE_MAP](#) regular
20 expression for the [job](#)
21 **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)
22 Comma-delimited list of ranks on this node within the specified namespace - referenced
23 using [PMIX_RANK_WILDCARD](#) .
24 **PMIX_LOCAL_PROCS** "pmix.lprocs" (pmix_proc_t array)
25 Array of [pmix_proc_t](#) of all processes on the specified node - referenced using
26 [PMIX_RANK_WILDCARD](#) .
27 **PMIX_LOCAL_CPUSSETS** "pmix.lcpus" (char*)
28 Colon-delimited cpusets of local peers within the specified namespace - referenced using
29 [PMIX_RANK_WILDCARD](#) .
30 **PMIX_PROC_URI** "pmix.puri" (char*)
31 URI containing contact information for a given process.
32 **PMIX_LOCALITY** "pmix.loc" (uint16_t)
33 Relative locality of the specified process to the requestor.
34 **PMIX_PARENT_ID** "pmix.parent" (pmix_proc_t)
35 Process identifier of the parent process of the calling process.
36 **PMIX_EXIT_CODE** "pmix.exit.code" (int)
37 Exit code returned when process terminated

38 **3.4.11 Information retrieval attributes**

39 The following attributes are used to specify the level of information (e.g., [session](#) , [job](#) , or
40 [application](#)) being requested where ambiguity may exist - see [5.1.5](#) for examples of their use.

41 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)

1 Return information about the specified session. If information about a session other than the
2 one containing the requesting process is desired, then the attribute array must contain a
3 [PMIX_SESSION_ID](#) attribute identifying the desired target.

4 **PMIX_JOB_INFO** "pmix.job.info" (bool)

5 Return information about the specified job or namespace. If information about a job or
6 namespace other than the one containing the requesting process is desired, then the attribute
7 array must contain a [PMIX_JOBID](#) or [PMIX_NAMESPACE](#) attribute identifying the desired
8 target. Similarly, if information is requested about a job or namespace in a session other than
9 the one containing the requesting process, then an attribute identifying the target session
10 must be provided.

11 **PMIX_APP_INFO** "pmix.app.info" (bool)

12 Return information about the specified application. If information about an application other
13 than the one containing the requesting process is desired, then the attribute array must
14 contain a [PMIX_APPNUM](#) attribute identifying the desired target. Similarly, if information
15 is requested about an application in a job or session other than the one containing the
16 requesting process, then attributes identifying the target job and/or session must be provided.

17 **PMIX_NODE_INFO** "pmix.node.info" (bool)

18 Return information about the specified node. If information about a node other than the one
19 containing the requesting process is desired, then the attribute array must contain either the
20 [PMIX_NODEID](#) or [PMIX_HOSTNAME](#) attribute identifying the desired target.

21 3.4.12 Information storage attributes

22 The following attributes are used to assemble information by its level (e.g., [session](#) , [job](#) , or
23 [application](#)) for storage where ambiguity may exist - see [11.1.3.1](#) for examples of their use.

24 **PMIX_SESSION_INFO_ARRAY** "pmix.ssn.arr" (pmix_data_array_t)

25 Provide an array of [pmix_info_t](#) containing session-level information. The
26 [PMIX_SESSION_ID](#) attribute is required to be included in the array.

27 **PMIX_JOB_INFO_ARRAY** "pmix.job.arr" (pmix_data_array_t)

28 Provide an array of [pmix_info_t](#) containing job-level information. The
29 [PMIX_SESSION_ID](#) attribute of the [session](#) containing the [job](#) is required to be
30 included in the array whenever the PMIx server library may host multiple sessions (e.g.,
31 when executing with a host RM daemon). As information is registered one job (aka
32 namespace) at a time via the [PMIx_server_register_namespace](#) API, there is no
33 requirement that the array contain either the [PMIX_NAMESPACE](#) or [PMIX_JOBID](#) attributes
34 when used in that context (though either or both of them may be included). At least one of
35 the job identifiers must be provided in all other contexts where the job being referenced is
36 ambiguous.

37 **PMIX_APP_INFO_ARRAY** "pmix.app.arr" (pmix_data_array_t)

38 Provide an array of [pmix_info_t](#) containing app-level information. The [PMIX_NAMESPACE](#)
39 or [PMIX_JOBID](#) attributes of the [job](#) containing the application, plus its [PMIX_APPNUM](#)
40 attribute, are must to be included in the array when the array is *not* included as part of a call
41 to [PMIx_server_register_namespace](#) - i.e., when the job containing the application is
42 ambiguous. The job identification is otherwise optional.

1 **PMIX_NODE_INFO_ARRAY** "pmix.node.arr" (pmix_data_array_t)

2 Provide an array of **pmix_info_t** containing node-level information. At a minimum,
3 either the **PMIX_NODEID** or **PMIX_HOSTNAME** attribute is required to be included in the
4 array, though both may be included.

5 Note that these assemblages can be used hierarchically:

- 6 • a **PMIX_JOB_INFO_ARRAY** might contain multiple **PMIX_APP_INFO_ARRAY** elements,
7 each describing values for a specific application within the job
- 8 • a **PMIX_JOB_INFO_ARRAY** could contain a **PMIX_NODE_INFO_ARRAY** for each node
9 hosting processes from that job, each array describing job-level values for that node
- 10 • a **PMIX_SESSION_INFO_ARRAY** might contain multiple **PMIX_JOB_INFO_ARRAY**
11 elements, each describing a job executing within the session. Each job array could, in turn,
12 contain both application and node arrays, thus providing a complete picture of the active
13 operations within the allocation

Advice to PMIx library implementers

14 PMIx implementations must be capable of properly parsing and storing any hierarchical depth of
15 information arrays. The resulting stored values are must to be accessible via both **PMIx_Get** and
16 **PMIx_Query_info_nb** APIs, assuming appropriate directives are provided by the caller.

3.4.13 Size information attributes

18 These attributes are used to describe the size of various dimensions of the PMIx universe - all are
19 referenced using **PMIX_RANK_WILDCARD**.

20 **PMIX_UNIV_SIZE** "pmix.univ.size" (uint32_t)

21 Number of allocated slots in a session - each slot may or may not be occupied by an
22 executing process. Note that this attribute is the equivalent to the combination of
23 **PMIX_SESSION_INFO_ARRAY** with the **PMIX_MAX_PROCS** entry in the array - it is
24 included in the Standard for historical reasons.

25 **PMIX_JOB_SIZE** "pmix.job.size" (uint32_t)

26 Total number of processes in this job across all contained applications. Note that this value
27 can be different from **PMIX_MAX_PROCS**. For example, users may choose to subdivide an
28 allocation (running several jobs in parallel within it), and dynamic programming models may
29 support adding and removing processes from a running **job** on-the-fly. In the latter case,
30 PMIx events must be used to notify processes within the job that the job size has changed.

31 **PMIX_JOB_NUM_APPS** "pmix.job.napps" (uint32_t)

32 Number of applications in this job.

33 **PMIX_APP_SIZE** "pmix.app.size" (uint32_t)

34 Number of processes in this application.

35 **PMIX_LOCAL_SIZE** "pmix.local.size" (uint32_t)

1 Number of processes in this job or application on this node.

2 **PMIX_NODE_SIZE** "pmix.node.size" (uint32_t)

3 Number of processes across all jobs on this node.

4 **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)

5 Maximum number of processes that can be executed in this context (session, namespace,
6 application, or node). Typically, this is a constraint imposed by a scheduler or by user
7 settings in a hostfile or other resource description.

8 **PMIX_NUM_SLOTS** "pmix.num.slots" (uint32_t)

9 Number of slots allocated in this context (session, namespace, application, or node). Note
10 that this attribute is the equivalent to **PMIX_MAX_PROCS** used in the corresponding
11 context - it is included in the Standard for historical reasons.

12 **PMIX_NUM_NODES** "pmix.num.nodes" (uint32_t)

13 Number of nodes in this session, or that are currently executing processes from the
14 associated namespace or application.

15 **3.4.14 Memory information attributes**

16 These attributes are used to describe memory available and used in the system - all are referenced
17 using **PMIX_RANK_WILDCARD**.

18 **PMIX_AVAIL_PHYS_MEMORY** "pmix.pmem" (uint64_t)

19 Total available physical memory on this node.

20 **PMIX_DAEMON_MEMORY** "pmix.dmn.mem" (float)

21 Megabytes of memory currently used by the RM daemon.

22 **PMIX_CLIENT_AVG_MEMORY** "pmix.cl.mem.avg" (float)

23 Average Megabytes of memory used by client processes.

24 **3.4.15 Topology information attributes**

25 These attributes are used to describe topology information in the PMIx universe - all are referenced
26 using **PMIX_RANK_WILDCARD** except where noted.

27 **PMIX_LOCAL_TOPO** "pmix.ltopo" (char*)

28 eXtensible Markup Language (XML) representation of local node topology.

29 **PMIX_TOPOLOGY** "pmix.topo" (hwloc_topology_t)

30 Pointer to the PMIx client's internal hwloc topology object.

31 **PMIX_TOPOLOGY_XML** "pmix.topo.xml" (char*)

32 XML-based description of topology

33 **PMIX_TOPOLOGY_FILE** "pmix.topo.file" (char*)

34 Full path to file containing XML topology description

35 **PMIX_TOPOLOGY_SIGNATURE** "pmix.toposig" (char*)

36 Topology signature string.

37 **PMIX_LOCALITY_STRING** "pmix.locstr" (char*)

String describing a process's bound location - referenced using the process's rank. The string is of the form:

NM%*s*:SK%*s*:L3%*s*:L2%*s*:L1%*s*:CR%*s*:HT%*s*

Where the *%s* is replaced with an integer index or inclusive range for hwloc. **NM** identifies the numa node(s). **SK** identifies the socket(s). **L3** identifies the L3 cache(s). **L2** identifies the L2 cache(s). **L1** identifies the L1 cache(s). **CR** identifies the cores(s). **HT** identifies the hardware thread(s). If your architecture does not have the specified hardware designation then it can be omitted from the signature.

For example: **NM0:SK0:L30-4:L20-4:L10-4:CR0-4:HT0-39**.

This means numa node 0, socket 0, L3 caches 0, 1, 2, 3, 4, L2 caches 0-4, L1 caches 0-4, cores 0, 1, 2, 3, 4, and hardware threads 0-39.

PMIX_HWLOC_SHMEM_ADDR "pmix.hwlocaddr" (**size_t**)

Address of the HWLOC shared memory segment.

PMIX_HWLOC_SHMEM_SIZE "pmix.hwlocsize" (**size_t**)

Size of the HWLOC shared memory segment.

PMIX_HWLOC_SHMEM_FILE "pmix.hwlocfile" (**char***)

Path to the HWLOC shared memory file.

PMIX_HWLOC_XML_V1 "pmix.hwlocxml1" (**char***)

XML representation of local topology using HWLOC's v1.x format.

PMIX_HWLOC_XML_V2 "pmix.hwlocxml2" (**char***)

XML representation of local topology using HWLOC's v2.x format.

PMIX_HWLOC_SHARE_TOPO "pmix.hwlocsh" (**bool**)

Share the HWLOC topology via shared memory

PMIX_HWLOC_HOLE_KIND "pmix.hwlocholek" (**char***)

Kind of VM "hole" HWLOC should use for shared memory

3.4.16 Request-related attributes

These attributes are used to influence the behavior of various PMIx operations - they do not represent values accessed using the [PMIx_Get](#) API.

PMIX_COLLECT_DATA "pmix.collect" (**bool**)

Collect data and return it at the end of the operation.

PMIX_TIMEOUT "pmix.timeout" (**int**)

Time in seconds before the specified operation should time out (*0* indicating infinite) in error. The timeout parameter can help avoid "hangs" due to programming errors that prevent the target process from ever exposing its data.

PMIX_IMMEDIATE "pmix.immediate" (**bool**)

Specified operation should immediately return an error from the PMIx server if the requested data cannot be found - do not request it from the host RM.

PMIX_WAIT "pmix.wait" (**int**)

Caller requests that the PMIx server wait until at least the specified number of values are found (*0* indicates all and is the default).

PMIX_COLLECTIVE_ALGO "pmix.calgo" (**char***)

1 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
2 impose any requirements on a host environment's collective algorithms. Thus, the
3 acceptable values for this attribute will be environment-dependent - users are encouraged to
4 check their host environment for supported values.

5 **PMIX_COLLECTIVE_ALGO_REQD** "pmix.calreqd" (bool)

6 If **true**, indicates that the requested choice of algorithm is mandatory.

7 **PMIX_NOTIFY_COMPLETION** "pmix.notecomp" (bool)

8 Notify the parent process upon termination of child job.

9 **PMIX_RANGE** "pmix.range" (pmix_data_range_t)

10 Value for calls to publish/lookup/unpublish or for monitoring event notifications.

11 **PMIX_PERSISTENCE** "pmix.persist" (pmix_persistence_t)

12 Value for calls to [PMIx_Publish](#).

13 **PMIX_DATA_SCOPE** "pmix.scope" (pmix_scope_t)

14 Scope of the data to be found in a [PMIx_Get](#) call.

15 **PMIX_OPTIONAL** "pmix.optional" (bool)

16 Look only in the client's local data store for the requested value - do not request data from
17 the PMIx server if not found.

18 **PMIX_EMBED_BARRIER** "pmix.embed.barrier" (bool)

19 Execute a blocking fence operation before executing the specified operation. For example,
20 [PMIx_Finalize](#) does not include an internal barrier operation by default. This attribute
21 would direct [PMIx_Finalize](#) to execute a barrier as part of the finalize operation.

22 **PMIX_JOB_TERM_STATUS** "pmix.job.term.status" (pmix_status_t)

23 Status returned by job upon its termination. The status will be communicated as part of a
24 PMIx event payload provided by the host environment upon termination of a job. Note that
25 generation of the [PMIX_ERR_JOB_TERMINATED](#) event is optional and host environments
26 may choose to provide it only upon request.

27 **PMIX_PROC_STATE_STATUS** "pmix.proc.state" (pmix_proc_state_t)

28 State of the specified process as of the last report - may not be the actual current state based
29 on update rate.

30 **PMIX_PROC_TERM_STATUS** "pmix.proc.term.status" (pmix_status_t)

31 Status returned by a process upon its termination. The status will be communicated as part
32 of a PMIx event payload provided by the host environment upon termination of a process.
33 Note that generation of the [PMIX_PROC_TERMINATED](#) event is optional and host
34 environments may choose to provide it only upon request.

35 3.4.17 Server-to-PMIx library attributes

36 Attributes used by the host environment to pass data to its PMIx server library. The data will then
37 be parsed and provided to the local PMIx clients. These attributes are all referenced using
38 [PMIX_RANK_WILDCARD](#) except where noted.

39 **PMIX_REGISTER_NODATA** "pmix.reg.nodata" (bool)

40 Registration is for this namespace only, do not copy job data - this attribute is not accessed
41 using the [PMIx_Get](#)

1 **PMIX_PROC_DATA** "pmix.pdata" (**pmix_data_array_t**)
2 Array of process data. Starts with rank, then contains more data.
3 **PMIX_NODE_MAP** "pmix.nmap" (**char***)
4 Regular expression of nodes - see 11.1.3.1 for an explanation of its generation.
5 **PMIX_PROC_MAP** "pmix.pmap" (**char***)
6 Regular expression describing processes on each node - see 11.1.3.1 for an explanation of its
7 generation.
8 **PMIX_ANL_MAP** "pmix.anlmap" (**char***)
9 Process mapping in Argonne National Laboratory's PMI-1/PMI-2 notation.
10 **PMIX_APP_MAP_TYPE** "pmix.apmap.type" (**char***)
11 Type of mapping used to layout the application (e.g., **cyclic**).
12 **PMIX_APP_MAP_REGEX** "pmix.apmap.regex" (**char***)
13 Regular expression describing the result of the process mapping.

14 3.4.18 Server-to-Client attributes

15 Attributes used internally to communicate data from the PMIx server to the PMIx client - they do
16 not represent values accessed using the **PMIx_Get** API.

17 **PMIX_PROC_BLOB** "pmix.pblob" (**pmix_byte_object_t**)
18 Packed blob of process data.
19 **PMIX_MAP_BLOB** "pmix.mblob" (**pmix_byte_object_t**)
20 Packed blob of process location.

21 3.4.19 Event handler registration and notification attributes

22 Attributes to support event registration and notification - they are values passed to the event
23 registration and notification APIs and therefore are not accessed using the **PMIx_Get** API.

24 **PMIX_EVENT_HDLR_NAME** "pmix.evname" (**char***)
25 String name identifying this handler.
26 **PMIX_EVENT_HDLR_FIRST** "pmix.evfirst" (**bool**)
27 Invoke this event handler before any other handlers.
28 **PMIX_EVENT_HDLR_LAST** "pmix.evlast" (**bool**)
29 Invoke this event handler after all other handlers have been called.
30 **PMIX_EVENT_HDLR_FIRST_IN_CATEGORY** "pmix.evfirstcat" (**bool**)
31 Invoke this event handler before any other handlers in this category.
32 **PMIX_EVENT_HDLR_LAST_IN_CATEGORY** "pmix.evlastcat" (**bool**)
33 Invoke this event handler after all other handlers in this category have been called.
34 **PMIX_EVENT_HDLR_BEFORE** "pmix.evbefore" (**char***)
35 Put this event handler immediately before the one specified in the (**char***) value.
36 **PMIX_EVENT_HDLR_AFTER** "pmix.evafter" (**char***)
37 Put this event handler immediately after the one specified in the (**char***) value.
38 **PMIX_EVENT_HDLR_PREPEND** "pmix.evprepend" (**bool**)
39 Prepend this handler to the precedence list within its category.

1 **PMIX_EVENT_HDLR_APPEND** "pmix.evappend" (bool)
 2 Append this handler to the precedence list within its category.
 3 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrangle" (pmix_data_array_t*)
 4 Array of **pmix_proc_t** defining range of event notification.
 5 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)
 6 The single process that was affected.
 7 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (pmix_data_array_t*)
 8 Array of **pmix_proc_t** defining affected processes.
 9 **PMIX_EVENT_NON_DEFAULT** "pmix.evnondf" (bool)
 10 Event is not to be delivered to default event handlers.
 11 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (void *)
 12 Object to be returned whenever the registered callback function **cbfunc** is invoked. The
 13 object will only be returned to the process that registered it.
 14 **PMIX_EVENT_DO_NOT_CACHE** "pmix.evnocache" (bool)
 15 Instruct the PMIx server not to cache the event.
 16 **PMIX_EVENT_SILENT_TERMINATION** "pmix.evsilentterm" (bool)
 17 Do not generate an event when this job normally terminates.
 18 **PMIX_EVENT_PROXY** "pmix.evproxy" (pmix_proc_t*)
 19 PMIx server that sourced the event
 20 **PMIX_EVENT_TEXT_MESSAGE** "pmix.evtext" (char*)
 21 Text message suitable for output by recipient - e.g., describing the cause of the event

22 3.4.20 Fault tolerance attributes

23 Attributes to support fault tolerance behaviors - they are values passed to the event notification API
 24 and therefore are not accessed using the **PMIx_Get** API.

25 **PMIX_EVENT_TERMINATE_SESSION** "pmix.evterm.sess" (bool)
 26 The RM intends to terminate this session.
 27 **PMIX_EVENT_TERMINATE_JOB** "pmix.evterm.job" (bool)
 28 The RM intends to terminate this job.
 29 **PMIX_EVENT_TERMINATE_NODE** "pmix.evterm.node" (bool)
 30 The RM intends to terminate all processes on this node.
 31 **PMIX_EVENT_TERMINATE_PROC** "pmix.evterm.proc" (bool)
 32 The RM intends to terminate just this process.
 33 **PMIX_EVENT_ACTION_TIMEOUT** "pmix.evtimeout" (int)
 34 The time in seconds before the RM will execute error response.
 35 **PMIX_EVENT_NO_TERMINATION** "pmix.evnoterm" (bool)
 36 Indicates that the handler has satisfactorily handled the event and believes termination of the
 37 application is not required.
 38 **PMIX_EVENT_WANT_TERMINATION** "pmix.evterm" (bool)
 39 Indicates that the handler has determined that the application should be terminated

1 3.4.21 Spawn attributes

2 Attributes used to describe **PMIx_Spawn** behavior - they are values passed to the **PMIx_Spawn**
3 API and therefore are not accessed using the **PMIx_Get** API when used in that context. However,
4 some of the attributes defined in this section can be provided by the host environment for other
5 purposes - e.g., the environment might provide the **PMIX_MAPPER** attribute in the job-related
6 information so that an application can use **PMIx_Get** to discover the layout algorithm used for
7 determining process locations. Multi-use attributes and their respective access reference rank are
8 denoted below.

9 **PMIX_PERSONALITY** "pmix.pers" (char*)

10 Name of personality to use.

11 **PMIX_HOST** "pmix.host" (char*)

12 Comma-delimited list of hosts to use for spawned processes.

13 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

14 Hostfile to use for spawned processes.

15 **PMIX_ADD_HOST** "pmix.addhost" (char*)

16 Comma-delimited list of hosts to add to the allocation.

17 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

18 Hostfile listing hosts to add to existing allocation.

19 **PMIX_PREFIX** "pmix.prefix" (char*)

20 Prefix to use for starting spawned processes.

21 **PMIX_WDIR** "pmix.wdir" (char*)

22 Working directory for spawned processes.

23 **PMIX_MAPPER** "pmix.mapper" (char*)

24 Mapping mechanism to use for placing spawned processes - when accessed using
25 **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping
26 mechanism used for the provided namespace.

27 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)

28 Display process mapping upon spawn.

29 **PMIX_PPR** "pmix.ppr" (char*)

30 Number of processes to spawn on each identified resource.

31 **PMIX_MAPBY** "pmix.mapby" (char*)

32 Process mapping policy - when accessed using **PMIx_Get**, use the
33 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
34 provided namespace

35 **PMIX_RANKBY** "pmix.rankby" (char*)

36 Process ranking policy - when accessed using **PMIx_Get**, use the
37 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the
38 provided namespace

39 **PMIX_BINDTO** "pmix.bindto" (char*)

40 Process binding policy - when accessed using **PMIx_Get**, use the
41 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
42 provided namespace

1 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)
 2 Preload binaries onto nodes.

3 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)
 4 Comma-delimited list of files to pre-position on nodes.

5 **PMIX_NON_PMI** "pmix.nonpmi" (bool)
 6 Spawned processes will not call **PMIx_Init** .

7 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)
 8 Spawned process rank that is to receive **stdin**.

9 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (bool)
 10 Forward this process's **stdin** to the designated process.

11 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)
 12 Forward **stdout** from spawned processes to this process.

13 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)
 14 Forward **stderr** from spawned processes to this process.

15 **PMIX_FWD_STDDIAG** "pmix.fwd.stddiag" (bool)
 16 If a diagnostic channel exists, forward any output on it from the spawned processes to this
 17 process (typically used by a tool)

18 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)
 19 Spawned application consists of debugger daemons.

20 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)
 21 Designated application is to be spawned as a disconnected job. Meaning that it is not part of
 22 the "comm_world" of the parent process.

23 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)
 24 Set the application's current working directory to the session working directory assigned by
 25 the RM - when accessed using **PMIx_Get** , use the **PMIX_RANK_WILDCARD** value for
 26 the rank to discover the session working directory assigned to the provided namespace

27 **PMIX_TAG_OUTPUT** "pmix.tagout" (bool)
 28 Tag application output with the identity of the source process.

29 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
 30 Timestamp output from applications.

31 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)
 32 Merge **stdout** and **stderr** streams from application processes.

33 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)
 34 Output application output to the specified file.

35 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)
 36 Mark the **argv** with the rank of the process.

37 **PMIX_CPUS_PER_PROC** "pmix.cpusperproc" (uint32_t)
 38 Number of cpus to assign to each rank - when accessed using **PMIx_Get** , use the
 39 **PMIX_RANK_WILDCARD** value for the rank to discover the cpus/process assigned to the
 40 provided namespace

41 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)
 42 Do not place processes on the head node.

43 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)

1 Do not oversubscribe the cpus.

2 **PMIX_REPORT_BINDINGS** "pmix.repbind" (bool)

3 Report bindings of the individual processes.

4 **PMIX_CPU_LIST** "pmix.cpulist" (char*)

5 List of cpus to use for this job - when accessed using **PMIx_Get**, use the

6 **PMIX_RANK_WILDCARD** value for the rank to discover the cpu list used for the provided
7 namespace

8 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)

9 Application supports recoverable operations.

10 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)

11 Application is continuous, all failed processes should be immediately restarted.

12 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)

13 Maximum number of times to restart a job - when accessed using **PMIx_Get**, use the

14 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
15 namespace

16 **PMIX_SPAWN_TOOL** "pmix.spwn.tool" (bool)

17 Indicate that the job being spawned is a tool

18 3.4.22 Query attributes

19 Attributes used to describe **PMIx_Query_info_nb** behavior - these are values passed to the
20 **PMIx_Query_info_nb** API and therefore are not passed to the **PMIx_Get** API.

21 **PMIX_QUERY_REFRESH_CACHE** "pmix.qry.rfsh" (bool)

22 Retrieve updated information from server.

23 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (char*)

24 Request a comma-delimited list of active namespaces.

25 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)

26 Status of a specified, currently executing job.

27 **PMIX_QUERY_QUEUE_LIST** "pmix.qry.qlst" (char*)

28 Request a comma-delimited list of scheduler queues.

29 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (TBD)

30 Status of a specified scheduler queue.

31 **PMIX_QUERY_PROC_TABLE** "pmix.qry.ptable" (char*)

32 Input namespace of the job whose information is being requested returns (
33 **pmix_data_array_t**) an array of **pmix_proc_info_t**.

34 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lptable" (char*)

35 Input namespace of the job whose information is being requested returns (
36 **pmix_data_array_t**) an array of **pmix_proc_info_t** for processes in job on same
37 node.

38 **PMIX_QUERY_AUTHORIZATIONS** "pmix.qry.auths" (bool)

39 Return operations the PMIx tool is authorized to perform.

40 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)

41 Return a comma-delimited list of supported spawn attributes.

1 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)
 2 Return a comma-delimited list of supported debug attributes.
 3 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)
 4 Return information on memory usage for the processes indicated in the qualifiers.
 5 **PMIX_QUERY_LOCAL_ONLY** "pmix.qry.local" (bool)
 6 Constrain the query to local information only.
 7 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
 8 Report only average values for sampled information.
 9 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
 10 Report minimum and maximum values.
 11 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 12 String identifier of the allocation whose status is being requested.
 13 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
 14 Query number of seconds (uint32_t) remaining in allocation for the specified namespace.
 15 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)
 16 Query list of supported attributes for specified APIs
 17 **PMIX_QUERY_NUM_PSETS** "pmix.qry.psetnum" (size_t)
 18 Return the number of psets defined in the specified range (defaults to session).
 19 **PMIX_QUERY_PSET_NAMES** "pmix.qry.psets" (char*)
 20 Return a comma-delimited list of the names of the psets defined in the specified range
 21 (defaults to session).

22 3.4.23 Log attributes

23 Attributes used to describe **PMIx_Log_nb** behavior - these are values passed to the
 24 **PMIx_Log_nb** API and therefore are not accessed using the **PMIx_Get** API.
 25 **PMIX_LOG_SOURCE** "pmix.log.source" (pmix_proc_t*)
 26 ID of source of the log request
 27 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)
 28 Log string to **stderr**.
 29 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)
 30 Log string to **stdout**.
 31 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)
 32 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
 33 otherwise to local syslog
 34 **PMIX_LOG_LOCAL_SYSLOG** "pmix.log.lsys" (char*)
 35 Log data to local syslog. Defaults to **ERROR** priority.
 36 **PMIX_LOG_GLOBAL_SYSLOG** "pmix.log.gsys" (char*)
 37 Forward data to system "gateway" and log msg to that syslog Defaults to **ERROR** priority.
 38 **PMIX_LOG_SYSLOG_PRI** "pmix.log.syspri" (int)
 39 Syslog priority level
 40 **PMIX_LOG_TIMESTAMP** "pmix.log.tstmp" (time_t)
 41 Timestamp for log report

1 **PMIX_LOG_GENERATE_TIMESTAMP** "pmix.log.gtstmp" (bool)
 2 Generate timestamp for log
 3 **PMIX_LOG_TAG_OUTPUT** "pmix.log.tag" (bool)
 4 Label the output stream with the channel name (e.g., "stdout")
 5 **PMIX_LOG_TIMESTAMP_OUTPUT** "pmix.log.tsout" (bool)
 6 Print timestamp in output string
 7 **PMIX_LOG_XML_OUTPUT** "pmix.log.xml" (bool)
 8 Print the output stream in XML format
 9 **PMIX_LOG_ONCE** "pmix.log.once" (bool)
 10 Only log this once with whichever channel can first support it, taking the channels in priority
 11 order
 12 **PMIX_LOG_MSG** "pmix.log.msg" (pmix_byte_object_t)
 13 Message blob to be sent somewhere.
 14 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)
 15 Log via email based on [pmix_info_t](#) containing directives.
 16 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)
 17 Comma-delimited list of email addresses that are to receive the message.
 18 **PMIX_LOG_EMAIL_SENDER_ADDR** "pmix.log.emfaddr" (char*)
 19 Return email address of sender
 20 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)
 21 Subject line for email.
 22 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)
 23 Message to be included in email.
 24 **PMIX_LOG_EMAIL_SERVER** "pmix.log.esrvr" (char*)
 25 Hostname (or IP address) of estmp server
 26 **PMIX_LOG_EMAIL_SRVR_PORT** "pmix.log.esrvrprt" (int32_t)
 27 Port the email server is listening to
 28 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)
 29 Store the log data in a global data store (e.g., database)
 30 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)
 31 Log the provided information to the host environment's job record

32 **3.4.24 Debugger attributes**

33 Attributes used to assist debuggers - these are values that can be passed to the [PMIx_Spawn](#) or
 34 [PMIx_Init](#) APIs. Some may be accessed using the [PMIx_Get](#) API with the
 35 [PMIX_RANK_WILDCARD](#) rank.

36 **PMIX_DEBUG_STOP_ON_EXEC** "pmix.dbg.exec" (bool)
 37 Passed to [PMIx_Spawn](#) to indicate that the specified application is being spawned under
 38 debugger, and that the launcher is to pause the resulting application processes on first
 39 instruction for debugger attach.
 40 **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)

1 Passed to `PMIx_Spawn` to indicate that the specified application is being spawned under
2 debugger, and that the `PMIx` client library is to pause the resulting application processes
3 during `PMIx_Init` until debugger attach and release.

4 `PMIX_DEBUG_WAIT_FOR_NOTIFY` "pmix.dbg.notify" (bool)

5 Passed to `PMIx_Spawn` to indicate that the specified application is being spawned under
6 debugger, and that the resulting application processes are to pause at some
7 application-determined location until debugger attach and release.

8 `PMIX_DEBUG_JOB` "pmix.dbg.job" (char*)

9 Namespace of the job to be debugged - provided to the debugger upon launch.

10 `PMIX_DEBUG_WAITING_FOR_NOTIFY` "pmix.dbg.waiting" (bool)

11 Job to be debugged is waiting for a release - this is not a value accessed using the
12 `PMIx_Get` API.

13 `PMIX_DEBUG_JOB_DIRECTIVES` "pmix.dbg.jdirs" (pmix_data_array_t*)

14 Array of job-level directives

15 `PMIX_DEBUG_APP_DIRECTIVES` "pmix.dbg.adirs" (pmix_data_array_t*)

16 Array of app-level directives

17 3.4.25 Resource manager attributes

18 Attributes used to describe the RM - these are values assigned by the host environment and accessed
19 using the `PMIx_Get` API. The value of the provided namespace is unimportant but should be
20 given as the namespace of the requesting process and a rank of `PMIX_RANK_WILDCARD` used to
21 indicate that the information will be found with the job-level information.

22 `PMIX_RM_NAME` "pmix.rm.name" (char*)

23 String name of the RM.

24 `PMIX_RM_VERSION` "pmix.rm.version" (char*)

25 RM version string.

26 3.4.26 Environment variable attributes

27 Attributes used to adjust environment variables - these are values passed to the `PMIx_Spawn` API
28 and are not accessed using the `PMIx_Get` API.

29 `PMIX_SET_ENVAR` "pmix.envar.set" (pmix_envar_t*)

30 Set the envar to the given value, overwriting any pre-existing one

31 `PMIX_UNSET_ENVAR` "pmix.envar.unset" (char*)

32 Unset the environment variable specified in the string.

33 `PMIX_ADD_ENVAR` "pmix.envar.add" (pmix_envar_t*)

34 Add the environment variable, but do not overwrite any pre-existing one

35 `PMIX_PREPEND_ENVAR` "pmix.envar.prepnd" (pmix_envar_t*)

36 Prepend the given value to the specified environmental value using the given separator
37 character, creating the variable if it doesn't already exist

38 `PMIX_APPEND_ENVAR` "pmix.envar.appnd" (pmix_envar_t*)

39 Append the given value to the specified environmental value using the given separator
40 character, creating the variable if it doesn't already exist

1 3.4.27 Job Allocation attributes

2 Attributes used to describe the job allocation - these are values passed to the
3 [PMIx_Allocation_request_nb](#) API and are not accessed using the [PMIx_Get](#) API

4 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)

5 Provide a string identifier for this allocation request which can later be used to query status
6 of the request.

7 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)

8 The number of nodes.

9 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)

10 Regular expression of the specific nodes.

11 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)

12 Number of cpus.

13 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)

14 Regular expression of the number of cpus for each node.

15 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)

16 Regular expression of the specific cpus indicating the cpus involved.

17 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)

18 Number of Megabytes.

19 **PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)

20 Array of [pmix_info_t](#) describing requested network resources. This must include at
21 least: [PMIX_ALLOC_NETWORK_ID](#), [PMIX_ALLOC_NETWORK_TYPE](#), and
22 [PMIX_ALLOC_NETWORK_ENDPTS](#), plus whatever other descriptors are desired.

23 **PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)

24 The key to be used when accessing this requested network allocation. The allocation will be
25 returned/stored as a [pmix_data_array_t](#) of [pmix_info_t](#) indexed by this key and
26 containing at least one entry with the same key and the allocated resource description. The
27 type of the included value depends upon the network support. For example, a TCP allocation
28 might consist of a comma-delimited string of socket ranges such as

29 "32000-32100,33005,38123-38146". Additional entries will consist of any provided

30 resource request directives, along with their assigned values. Examples include:

31 [PMIX_ALLOC_NETWORK_TYPE](#) - the type of resources provided;

32 [PMIX_ALLOC_NETWORK_PLANE](#) - if applicable, what plane the resources were assigned

33 from; [PMIX_ALLOC_NETWORK_QOS](#) - the assigned QoS; [PMIX_ALLOC_BANDWIDTH](#) -

34 the allocated bandwidth; [PMIX_ALLOC_NETWORK_SEC_KEY](#) - a security key for the

35 requested network allocation. NOTE: the assigned values may differ from those requested,

36 especially if [PMIX_INFO_REQD](#) was not set in the request.

37 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)

38 Mbits/sec.

39 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)

40 Quality of service level.

41 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)

42 Time in seconds.

```

1 PMIX_ALLOC_NETWORK_TYPE "pmix.alloc.nettype" (char*)
2     Type of desired transport (e.g., "tcp", "udp")
3 PMIX_ALLOC_NETWORK_PLANE "pmix.alloc.netplane" (char*)
4     ID string for the NIC (aka plane) to be used for this allocation (e.g., CIDR for Ethernet)
5 PMIX_ALLOC_NETWORK_ENDPTS "pmix.alloc.endpts" (size_t)
6     Number of endpoints to allocate per process
7 PMIX_ALLOC_NETWORK_ENDPTS_NODE "pmix.alloc.endpts.nd" (size_t)
8     Number of endpoints to allocate per node
9 PMIX_ALLOC_NETWORK_SEC_KEY "pmix.alloc.nsec" (pmix_byte_object_t)
10    Network security key

```

11 3.4.28 Job control attributes

12 Attributes used to request control operations on an executing application - these are values passed
13 to the [PMIx_Job_control_nb](#) API and are not accessed using the [PMIx_Get](#) API.

```

14 PMIX_JOB_CTRL_ID "pmix.jctrl.id" (char*)
15     Provide a string identifier for this request. The user can provide an identifier for the
16     requested operation, thus allowing them to later request status of the operation or to
17     terminate it. The host, therefore, shall track it with the request for future reference.
18 PMIX_JOB_CTRL_PAUSE "pmix.jctrl.pause" (bool)
19     Pause the specified processes.
20 PMIX_JOB_CTRL_RESUME "pmix.jctrl.resume" (bool)
21     Resume ("un-pause") the specified processes.
22 PMIX_JOB_CTRL_CANCEL "pmix.jctrl.cancel" (char*)
23     Cancel the specified request - the provided request ID must match the
24     PMIX\_JOB\_CTRL\_ID provided to a previous call to PMIx\_Job\_control . An ID of
25     NULL implies cancel all requests from this requestor.
26 PMIX_JOB_CTRL_KILL "pmix.jctrl.kill" (bool)
27     Forcibly terminate the specified processes and cleanup.
28 PMIX_JOB_CTRL_RESTART "pmix.jctrl.restart" (char*)
29     Restart the specified processes using the given checkpoint ID.
30 PMIX_JOB_CTRL_CHECKPOINT "pmix.jctrl.ckpt" (char*)
31     Checkpoint the specified processes and assign the given ID to it.
32 PMIX_JOB_CTRL_CHECKPOINT_EVENT "pmix.jctrl.ckptev" (bool)
33     Use event notification to trigger a process checkpoint.
34 PMIX_JOB_CTRL_CHECKPOINT_SIGNAL "pmix.jctrl.ckptsig" (int)
35     Use the given signal to trigger a process checkpoint.
36 PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT "pmix.jctrl.ckptsig" (int)
37     Time in seconds to wait for a checkpoint to complete.
38 PMIX_JOB_CTRL_CHECKPOINT_METHOD
39     "pmix.jctrl.ckmethod" (pmix_data_array_t)
40     Array of pmix\_info\_t declaring each method and value supported by this application.
41 PMIX_JOB_CTRL_SIGNAL "pmix.jctrl.sig" (int)

```

1 Send given signal to specified processes.

2 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)

3 Regular expression identifying nodes that are to be provisioned.

4 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)

5 Name of the image that is to be provisioned.

6 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)

7 Indicate that the job can be pre-empted.

8 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)

9 Politely terminate the specified processes.

10 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)

11 Comma-delimited list of files to be removed upon process termination

12 **PMIX_REGISTER_CLEANUP_DIR** "pmix.reg.cleanupdir" (char*)

13 Comma-delimited list of directories to be removed upon process termination

14 **PMIX_CLEANUP_RECURSIVE** "pmix.clnup.recurse" (bool)

15 Recursively cleanup all subdirectories under the specified one(s)

16 **PMIX_CLEANUP_EMPTY** "pmix.clnup.empty" (bool)

17 Only remove empty subdirectories

18 **PMIX_CLEANUP_IGNORE** "pmix.clnup.ignore" (char*)

19 Comma-delimited list of filenames that are not to be removed

20 **PMIX_CLEANUP_LEAVE_TOPDIR** "pmix.clnup.lvtop" (bool)

21 When recursively cleaning subdirectories, do not remove the top-level directory (the one
22 given in the cleanup request)

23 3.4.29 Monitoring attributes

24 Attributes used to control monitoring of an executing application- these are values passed to the
25 [PMIx_Process_monitor_nb](#) API and are not accessed using the [PMIx_Get](#) API.

26 **PMIX_MONITOR_ID** "pmix.monitor.id" (char*)

27 Provide a string identifier for this request.

28 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)

29 Identifier to be canceled (NULL means cancel all monitoring for this process).

30 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)

31 The application desires to control the response to a monitoring event.

32 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)

33 Register to have the PMIx server monitor the requestor for heartbeats.

34 **PMIX_SEND_HEARTBEAT** "pmix.monitor.beat" (void)

35 Send heartbeat to local PMIx server.

36 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)

37 Time in seconds before declaring heartbeat missed.

38 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrops" (uint32_t)

39 Number of heartbeats that can be missed before generating the event.

40 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)

41 Register to monitor file for signs of life.

1 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)
 2 Monitor size of given file is growing to determine if the application is running.
 3 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)
 4 Monitor time since last access of given file to determine if the application is running.
 5 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
 6 Monitor time since last modified of given file to determine if the application is running.
 7 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
 8 Time in seconds between checking the file.
 9 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)
 10 Number of file checks that can be missed before generating the event.

11 3.4.30 Security attributes

12 *PMIx v3.0* Attributes for managing security credentials

13 **PMIX_CRED_TYPE** "pmix.sec.ctype" (char*)
 14 When passed in **PMIx_Get_credential**, a prioritized, comma-delimited list of desired
 15 credential types for use in environments where multiple authentication mechanisms may be
 16 available. When returned in a callback function, a string identifier of the credential type.
 17 **PMIX_CRYPTO_KEY** "pmix.sec.key" (pmix_byte_object_t)
 18 Blob containing crypto key

19 3.4.31 IO Forwarding attributes

20 *PMIx v3.0* Attributes used to control IO forwarding behavior

21 **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (uint32_t)
 22 The requested size of the server cache in bytes for each specified channel. By default, the
 23 server is allowed (but not required) to drop all bytes received beyond the max size.
 24 **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (bool)
 25 In an overflow situation, drop the oldest bytes to make room in the cache.
 26 **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (bool)
 27 In an overflow situation, drop any new bytes received until room becomes available in the
 28 cache (default).
 29 **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (uint32_t)
 30 Controls grouping of IO on the specified channel(s) to avoid being called every time a bit of
 31 IO arrives. The library will execute the callback whenever the specified number of bytes
 32 becomes available. Any remaining buffered data will be “flushed” upon call to deregister the
 33 respective channel.
 34 **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (uint32_t)
 35 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
 36 size, this prevents IO from being held indefinitely while waiting for another payload to arrive.
 37 **PMIX_IOF_COMPLETE** "pmix.iof.cmp" (bool)
 38 Indicates whether or not the specified IO channel has been closed by the source.
 39 **PMIX_IOF_TAG_OUTPUT** "pmix.iof.tag" (bool)

1 Tag output with the channel it comes from.
2 **PMIX_IOF_TIMESTAMP_OUTPUT** "pmix.iof.ts" (bool)
3 Timestamp output
4 **PMIX_IOF_XML_OUTPUT** "pmix.iof.xml" (bool)
5 Format output in XML

3.4.32 Application setup attributes

7 *PMIx v3.0* Attributes for controlling contents of application setup data
8 **PMIX_SETUP_APP_ENVARS** "pmix.setup.env" (bool)
9 Harvest and include relevant environmental variables
10 **PMIX_SETUP_APP_NONENVARS** "pmix.setup.nenv" (bool)
11 Include all relevant data other than environmental variables
12 **PMIX_SETUP_APP_ALL** "pmix.setup.all" (bool)
13 Include all relevant data

3.4.33 Attribute support level attributes

15 **PMIX_CLIENT_FUNCTIONS** "pmix.client.fns" (bool)
16 Request a list of functions supported by the PMIx client library
17 **PMIX_CLIENT_ATTRIBUTES** "pmix.client.attrs" (bool)
18 Request attributes supported by the PMIx client library
19 **PMIX_SERVER_FUNCTIONS** "pmix.srvr.fns" (bool)
20 Request a list of functions supported by the PMIx server library
21 **PMIX_SERVER_ATTRIBUTES** "pmix.srvr.attrs" (bool)
22 Request attributes supported by the PMIx server library
23 **PMIX_HOST_FUNCTIONS** "pmix.srvr.fns" (bool)
24 Request a list of functions supported by the host environment
25 **PMIX_HOST_ATTRIBUTES** "pmix.host.attrs" (bool)
26 Request attributes supported by the host environment
27 **PMIX_TOOL_FUNCTIONS** "pmix.tool.fns" (bool)
28 Request a list of functions supported by the PMIx tool library
29 **PMIX_TOOL_ATTRIBUTES** "pmix.setup.env" (bool)
30 Request attributes supported by the PMIx tool library functions

3.4.34 Descriptive attributes

32 **PMIX_MAX_VALUE** "pmix.descr.maxval" (varies)
33 Used in **pmix_regattr_t** to describe the maximum valid value for the associated
34 attribute.
35 **PMIX_MIN_VALUE** "pmix.descr.minval" (varies)
36 Used in **pmix_regattr_t** to describe the minimum valid value for the associated
37 attribute.
38 **PMIX_ENUM_VALUE** "pmix.descr.enum" (char*)
39 Used in **pmix_regattr_t** to describe accepted values for the associated attribute.
40 Numerical values shall be presented in a form convertible to the attribute's declared data
41 type. Named values (i.e., values defined by constant names via a typical C-language enum
42 declaration) must be provided as their numerical equivalent.

1 3.4.35 Process group attributes

2 *PMIx v4.0* Attributes for controlling the PMIx Group APIs

3 **PMIX_GROUP_ID** "pmix.grp.id" (char*)

4 User-provided group identifier

5 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)

6 This process is the leader of the group

7 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

8 Participation is optional - do not return an error if any of the specified processes terminate
9 without having joined. The default is false

10 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

11 Notify remaining members when another member terminates without first leaving the group.
12 The default is false

13 **PMIX_GROUP_INVITE_DECLINE** "pmix.grp.decline" (bool)

14 Notify the inviting process that this process does not wish to participate in the proposed
15 group The default is true

16 **PMIX_GROUP_MEMBERSHIP** "pmix.grp.mbrs" (pmix_data_array_t*)

17 Array of group member ID's

18 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

19 Requests that the RM assign a new context identifier to the newly created group. The
20 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
21 specified in the request. Thus, the value serves as a means of identifying the group within
22 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION** .

23 **PMIX_GROUP_CONTEXT_ID** "pmix.grp.ctxid" (size_t)

24 Context identifier assigned to the group by the host RM.

25 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)

26 Group operation only involves local processes. PMIx implementations are *required* to
27 automatically scan an array of group members for local vs remote processes - if only local
28 processes are detected, the implementation need not execute a global collective for the
29 operation unless a context ID has been requested from the host environment. This can result
30 in significant time savings. This attribute can be used to optimize the operation by indicating
31 whether or not only local processes are represented, thus allowing the implementation to
32 bypass the scan. The default is false

33 **PMIX_GROUP_ENDPT_DATA** "pmix.grp.endpt" (pmix_byte_object_t)

34 Data collected to be shared during group construction

35 3.5 Callback Functions

36 PMIx provides blocking and nonblocking versions of most APIs. In the nonblocking versions, a
37 callback is activated upon completion of the the operation. This section describes many of those
38 callbacks.

1 3.5.1 Release Callback Function

2 Summary

3 The `pmix_release_cbfunc_t` is used by the `pmix_modex_cbfunc_t` and
4 `pmix_info_cbfunc_t` operations to indicate that the callback data may be reclaimed/freed by
5 the caller.

6 Format

PMIx v1.0

▼ C ▼

```
7 typedef void (*pmix_release_cbfunc_t)
8     (void *cbdata)
```

▲ C ▲

9 INOUT cbdata

10 Callback data passed to original API call (memory reference)

11 Description

12 Since the data is “owned” by the host server, provide a callback function to notify the host server
13 that we are done with the data so it can be released.

14 3.5.2 Modex Callback Function

15 Summary

16 The `pmix_modex_cbfunc_t` is used by the `pmix_server_fencefn_t` and
17 `pmix_server_dmodex_req_fn_t` PMIx server operations to return modex business card
18 exchange (BCX) data.

PMIx v1.0

▼ C ▼

```
19 typedef void (*pmix_modex_cbfunc_t)
20     (pmix_status_t status,
21      const char *data, size_t ndata,
22      void *cbdata,
23      pmix_release_cbfunc_t release_fn,
24      void *release_cbdata)
```

▲ C ▲

25 IN status

26 Status associated with the operation (handle)

27 IN data

28 Data to be passed (pointer)

29 IN ndata

30 size of the data (`size_t`)

31 IN cbdata

32 Callback data passed to original API call (memory reference)

1 **IN** `release_fn`
2 Callback for releasing *data* (function pointer)
3 **IN** `release_cbdata`
4 Pointer to be passed to *release_fn* (memory reference)

5 **Description**

6 A callback function that is solely used by PMIx servers, and not clients, to return modex BCX data
7 in response to “fence” and “get” operations. The returned blob contains the data collected from
8 each server participating in the operation.

9 **3.5.3 Spawn Callback Function**

10 **Summary**

11 The `pmix_spawn_cbfunc_t` is used on the PMIx client side by `PMIx_Spawn_nb` and on
12 the PMIx server side by `pmix_server_spawn_fn_t`.

PMIx v1.0

```
▼ C ▼  
13 typedef void (*pmix_spawn_cbfunc_t)  
14     (pmix_status_t status,  
15      pmix_namespace_t nspace, void *cbdata);  
▲ C ▲
```

16 **IN** `status`
17 Status associated with the operation (handle)
18 **IN** `namespace`
19 Namespace string (`pmix_namespace_t`)
20 **IN** `cbdata`
21 Callback data passed to original API call (memory reference)

22 **Description**

23 The callback will be executed upon launch of the specified applications in `PMIx_Spawn_nb`, or
24 upon failure to launch any of them.

25 The *status* of the callback will indicate whether or not the spawn succeeded. The *namespace* of the
26 spawned processes will be returned, along with any provided callback data. Note that the returned
27 *namespace* value will not be protected by the PRI upon return from the callback function, so the
28 receiver must copy it if it needs to be retained.

1 3.5.4 Op Callback Function

2 Summary

3 The [pmix_op_cbfunc_t](#) is used by operations that simply return a status.

PMIx v1.0

```
▼ _____ C _____ ▼  
4 typedef void (*pmix_op_cbfunc_t)  
5     (pmix_status_t status, void *cbdata);  
▲ _____ C _____ ▲
```

6 **IN status**

7 Status associated with the operation (handle)

8 **IN cbdata**

9 Callback data passed to original API call (memory reference)

10 Description

11 Used by a wide range of PMIx API's including [PMIx_Fence_nb](#),
12 [pmix_server_client_connected_fn_t](#), [PMIx_server_register_namespace](#). This
13 callback function is used to return a status to an often nonblocking operation.

14 3.5.5 Lookup Callback Function

15 Summary

16 The [pmix_lookup_cbfunc_t](#) is used by [PMIx_Lookup_nb](#) to return data.

PMIx v1.0

```
▼ _____ C _____ ▼  
17 typedef void (*pmix_lookup_cbfunc_t)  
18     (pmix_status_t status,  
19     pmix_pdata_t data[], size_t ndata,  
20     void *cbdata);  
▲ _____ C _____ ▲
```

21 **IN status**

22 Status associated with the operation (handle)

23 **IN data**

24 Array of data returned ([pmix_pdata_t](#))

25 **IN ndata**

26 Number of elements in the *data* array (**size_t**)

27 **IN cbdata**

28 Callback data passed to original API call (memory reference)

Description

A callback function for calls to `PMIx_Lookup_nb`. The function will be called upon completion of the command with the *status* indicating the success or failure of the request. Any retrieved data will be returned in an array of `pmix_pdata_t` structs. The namespace and rank of the process that provided each data element is also returned.

Note that these structures will be released upon return from the callback function, so the receiver must copy/protect the data prior to returning if it needs to be retained.

3.5.6 Value Callback Function

Summary

The `pmix_value_cbfunc_t` is used by `PMIx_Get_nb` to return data.

PMIx v1.0

```
typedef void (*pmix_value_cbfunc_t)
    (pmix_status_t status,
     pmix_value_t *kv, void *cbdata);
```

IN status

Status associated with the operation (handle)

IN kv

Key/value pair representing the data (`pmix_value_t`)

IN cbdata

Callback data passed to original API call (memory reference)

Description

A callback function for calls to `PMIx_Get_nb`. The *status* indicates if the requested data was found or not. A pointer to the `pmix_value_t` structure containing the found data is returned. The pointer will be `NULL` if the requested data was not found.

3.5.7 Info Callback Function

Summary

The `pmix_info_cbfunc_t` is a general information callback used by various APIs.

PMIx v2.0

```
typedef void (*pmix_info_cbfunc_t)
    (pmix_status_t status,
     pmix_info_t info[], size_t ninfo,
     void *cbdata,
     pmix_release_cbfunc_t release_fn,
     void *release_cbdata);
```

C

1 **IN status**
2 Status associated with the operation (`pmix_status_t`)
3 **IN info**
4 Array of `pmix_info_t` returned by the operation (pointer)
5 **IN ninfo**
6 Number of elements in the *info* array (`size_t`)
7 **IN cbdata**
8 Callback data passed to original API call (memory reference)
9 **IN release_fn**
10 Function to be called when done with the *info* data (function pointer)
11 **IN release_cbdata**
12 Callback data to be passed to *release_fn* (memory reference)

Description

13 The *status* indicates if requested data was found or not. An array of `pmix_info_t` will contain
14 the key/value pairs.
15

16 3.5.8 Event Handler Registration Callback Function

17 The `pmix_evhdlr_reg_cbfunc_t` callback function.

Advice to users

18 The PMIx *ad hoc* v1.0 Standard defined an error handler registration callback function with a
19 compatible signature, but with a different type definition function name
20 (`pmix_errhandler_reg_cbfunc_t`). It was removed from the v2.0 Standard and is not included in this
21 document to avoid confusion.

PMIx v2.0

C

```
22 typedef void (*pmix_evhdlr_reg_cbfunc_t)  
23     (pmix_status_t status,  
24      size_t evhdlr_ref,  
25      void *cbdata)
```

C

26 **IN status**
27 Status indicates if the request was successful or not (`pmix_status_t`)
28 **IN evhdlr_ref**
29 Reference assigned to the event handler by PMIx — this reference * must be used to
30 deregister the err handler (`size_t`)
31 **IN cbdata**
32 Callback data passed to original API call (memory reference)

1 **Description**

2 Define a callback function for calls to [PMIx_Register_event_handler](#)

3 **3.5.9 Notification Handler Completion Callback Function**

4 **Summary**

5 The [pmix_event_notification_cbfnc_fn_t](#) is called by event handlers to indicate
6 completion of their operations.

PMIx v2.0

C

```
7 typedef void (*pmix_event_notification_cbfnc_fn_t)
8     (pmix_status_t status,
9      pmix_info_t *results, size_t nresults,
10     pmix_op_cbfnc_t cbfunc, void *thiscbdata,
11     void *notification_cbdata);
```

C

- 12 **IN status**
13 Status returned by the event handler's operation ([pmix_status_t](#))
- 14 **IN results**
15 Results from this event handler's operation on the event ([pmix_info_t](#))
- 16 **IN nresults**
17 Number of elements in the results array ([size_t](#))
- 18 **IN cbfunc**
19 [pmix_op_cbfnc_t](#) function to be executed when PMIx completes processing the
20 callback (function reference)
- 21 **IN thiscbdata**
22 Callback data that was passed in to the handler (memory reference)
- 23 **IN cbdata**
24 Callback data to be returned when PMIx executes cbfunc (memory reference)

25 **Description**

26 Define a callback by which an event handler can notify the PMIx library that it has completed its
27 response to the notification. The handler is *required* to execute this callback so the library can
28 determine if additional handlers need to be called. The handler shall return
29 [PMIX_EVENT_ACTION_COMPLETE](#) if no further action is required. The return status of each
30 event handler and any returned [pmix_info_t](#) structures will be added to the *results* array of
31 [pmix_info_t](#) passed to any subsequent event handlers to help guide their operation.

32 If non-NULL, the provided callback function will be called to allow the event handler to release the
33 provided info array and execute any other required cleanup operations.

1 3.5.10 Notification Function

2 Summary

3 The `pmix_notification_fn_t` is called by PMIx to deliver notification of an event.

Advice to users

4 The PMIx *ad hoc* v1.0 Standard defined an error notification function with an identical name, but
5 different signature than the v2.0 Standard described below. The *ad hoc* v1.0 version was removed
6 from the v2.0 Standard is not included in this document to avoid confusion.

PMIx v2.0

```
7 typedef void (*pmix_notification_fn_t)
8     (size_t evhdlr_registration_id,
9      pmix_status_t status,
10     const pmix_proc_t *source,
11     pmix_info_t info[], size_t ninfo,
12     pmix_info_t results[], size_t nresults,
13     pmix_event_notification_cbfnc_fn_t cbfunc,
14     void *cbdata);
```

15 **IN evhdlr_registration_id**

Registration number of the handler being called (**size_t**)

17 **IN status**

Status associated with the operation (**pmix_status_t**)

19 **IN source**

Identifier of the process that generated the event (**pmix_proc_t**). If the source is the SMS,
then the nspace will be empty and the rank will be PMIX_RANK_UNDEF

22 **IN info**

Information describing the event (**pmix_info_t**). This argument will be NULL if no
additional information was provided by the event generator.

25 **IN ninfo**

Number of elements in the info array (**size_t**)

27 **IN results**

Aggregated results from prior event handlers servicing this event (**pmix_info_t**). This
argument will be **NULL** if this is the first handler servicing the event, or if no prior handlers
provided results.

31 **IN nresults**

Number of elements in the results array (**size_t**)

33 **IN cbfunc**

pmix_event_notification_cbfnc_fn_t callback function to be executed upon
completion of the handler's operation and prior to handler return (function reference).

IN `cbdata`

Callback data to be passed to `cbfunc` (memory reference)

Description

Note that different RMs may provide differing levels of support for event notification to application processes. Thus, the *info* array may be **NULL** or may contain detailed information of the event. It is the responsibility of the application to parse any provided info array for defined key-values if it so desires.

Advice to users

Possible uses of the *info* array include:

- for the host RM to alert the process as to planned actions, such as aborting the session, in response to the reported event
- provide a timeout for alternative action to occur, such as for the application to request an alternate response to the event

For example, the RM might alert the application to the failure of a node that resulted in termination of several processes, and indicate that the overall session will be aborted unless the application requests an alternative behavior in the next 5 seconds. The application then has time to respond with a checkpoint request, or a request to recover from the failure by obtaining replacement nodes and restarting from some earlier checkpoint.

Support for these options is left to the discretion of the host RM. Info keys are included in the common definitions above but may be augmented by environment vendors.

Advice to PMIx server hosts

On the server side, the notification function is used to inform the PMIx server library's host of a detected event in the PMIx server library. Events generated by PMIx clients are communicated to the PMIx server library, but will be relayed to the host via the `pmix_server_notify_event_fn_t` function pointer, if provided.

3.5.11 Server Setup Application Callback Function

The `PMIx_server_setup_application` callback function.

Summary

Provide a function by which the resource manager can receive application-specific environmental variables and other setup data prior to launch of an application.

1
PMIx v2.0

Format

C

```
2 typedef void (*pmix_setup_application_cbfunc_t) (  
3         pmix_status_t status,  
4         pmix_info_t info[], size_t ninfo,  
5         void *provided_cbdata,  
6         pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

- 7 **IN status**
8 returned status of the request ([pmix_status_t](#))
- 9 **IN info**
10 Array of info structures (array of handles)
- 11 **IN ninfo**
12 Number of elements in the *info* array (integer)
- 13 **IN provided_cbdata**
14 Data originally passed to call to [PMIx_server_setup_application](#) (memory
15 reference)
- 16 **IN cbfunc**
17 [pmix_op_cbfunc_t](#) function to be called when processing completed (function reference)
- 18 **IN cbdata**
19 Data to be passed to the *cbfunc* callback function (memory reference)

Description

21 Define a function to be called by the PMIx server library for return of application-specific setup
22 data in response to a request from the host RM. The returned *info* array is owned by the PMIx
23 server library and will be free'd when the provided *cbfunc* is called.

24 3.5.12 Server Direct Modex Response Callback Function

25 The [PMIx_server_dmodex_request](#) callback function.

26 Summary

27 Provide a function by which the local PMIx server library can return connection and other data
28 posted by local application processes to the host resource manager.

29 Format

PMIx v1.0

C

```
30 typedef void (*pmix_dmodex_response_fn_t) (pmix_status_t status,  
31         char *data, size_t sz,  
32         void *cbdata);
```

C

1 **IN** **status**
2 Returned status of the request ([pmix_status_t](#))
3 **IN** **data**
4 Pointer to a data "blob" containing the requested information (handle)
5 **IN** **sz**
6 Number of bytes in the *data* blob (integer)
7 **IN** **cbdata**
8 Data passed into the initial call to [PMIx_server_dmodex_request](#) (memory reference)

Description

9
10 Define a function to be called by the PMIx server library for return of information posted by a local
11 application process (via [PMIx_Put](#) with subsequent [PMIx_Commit](#)) in response to a request
12 from the host RM. The returned *data* blob is owned by the PMIx server library and will be free'd
13 upon return from the function.

14 3.5.13 PMIx Client Connection Callback Function

15 Summary

16 Callback function for incoming connection request from a local client

17 Format

PMIx v1.0

C

```
18 typedef void (*pmix_connection_cbfunc_t) (  
19                                     int incoming_sd, void *cbdata)
```

C

20 **IN** **incoming_sd**
21 (integer)
22 **IN** **cbdata**
23 (memory reference)

24 Description

25 Callback function for incoming connection requests from local clients - only used by host
26 environments that wish to directly handle socket connection requests.

27 3.5.14 PMIx Tool Connection Callback Function

28 Summary

29 Callback function for incoming tool connections.

1
PMIx v2.0

Format

C

```
2 typedef void (*pmix_tool_connection_cbfunc_t) (  
3             pmix_status_t status,  
4             pmix_proc_t *proc, void *cbdata)
```

C

- 5 **IN** **status**
- 6 **pmix_status_t** value (handle)
- 7 **IN** **proc**
- 8 **pmix_proc_t** structure containing the identifier assigned to the tool (handle)
- 9 **IN** **cbdata**
- 10 Data to be passed (memory reference)

Description

11 Callback function for incoming tool connections. The host environment shall provide a
12 namespace/rank identifier for the connecting tool.
13

Advice to PMIx server hosts

14 It is assumed that **rank=0** will be the normal assignment, but allow for the future possibility of a
15 parallel set of tools connecting, and thus each process requiring a unique rank.

16 3.5.15 Credential callback function

17 Summary

18 Callback function to return a requested security credential

Format

C

```

typedef void (*pmix_credential_cbfunc_t) (
    pmix_status_t status,
    pmix_byte_object_t *credential,
    pmix_info_t info[], size_t ninfo,
    void *cbdata)

```

C

IN status`pmix_status_t` value (handle)**IN credential**`pmix_byte_object_t` structure containing the security credential (handle)**IN info**

Array of provided by the system to pass any additional information about the credential - e.g., the identity of the issuing agent. (handle)

IN ninfoNumber of elements in *info* (`size_t`)**IN cbdata**

Object passed in original request (memory reference)

Description

Define a callback function to return a requested security credential. Information provided by the issuing agent can subsequently be used by the application for a variety of purposes. Examples include:

- checking identified authorizations to determine what requests/operations are feasible as a means to steering **workflows**
- compare the credential type to that of the local SMS for compatibility

Advice to users

The credential is opaque and therefore understandable only by a service compatible with the issuer. The *info* array is owned by the PMIx library and is not to be released or altered by the receiving party.

28 3.5.16 Credential validation callback function**29 Summary**

30 Callback function for security credential validation

1
PMIx v3.0

Format

C

```
2 typedef void (*pmix_validation_cfunc_t) (  
3         pmix_status_t status,  
4         pmix_info_t info[], size_t ninfo,  
5         void *cbdata);
```

C

- 6 **IN status**
7 `pmix_status_t` value (handle)
- 8 **IN info**
9 Array of `pmix_info_t` provided by the system to pass any additional information about
10 the authentication - e.g., the effective userid and group id of the certificate holder, and any
11 related authorizations (handle)
- 12 **IN ninfo**
13 Number of elements in *info* (`size_t`)
- 14 **IN cbdata**
15 Object passed in original request (memory reference)

Description

16 Define a validation callback function to indicate if a provided credential is valid, and any
17 corresponding information regarding authorizations and other security matters.
18

Advice to users

19 The precise contents of the array will depend on the host environment and its associated security
20 system. At the minimum, it is expected (but not required) that the array will contain entries for the
21 `PMIX_USERID` and `PMIX_GRPID` of the client described in the credential. The *info* array is
22 owned by the PMIx library and is not to be released or altered by the receiving party.

23 3.5.17 IOF delivery function

24 Summary

25 Callback function for delivering forwarded IO to a process

1
PMIx v3.0

Format

C

```
2 typedef void (*pmix_iof_cbfunc_t) (  
3         size_t iofhdlr, pmix_iof_channel_t channel,  
4         pmix_proc_t *source, char *payload,  
5         pmix_info_t info[], size_t ninfo);
```

C

- 6 **IN** `iofhdlr`
7 Registration number of the handler being invoked (`size_t`)
- 8 **IN** `channel`
9 bitmask identifying the channel the data arrived on (`pmix_iof_channel_t`)
- 10 **IN** `source`
11 Pointer to a `pmix_proc_t` identifying the namespace/rank of the process that generated the
12 data (`char*`)
- 13 **IN** `payload`
14 Pointer to character array containing the data.
- 15 **IN** `info`
16 Array of `pmix_info_t` provided by the source containing metadata about the payload.
17 This could include `PMIX_IOF_COMPLETE` (handle)
- 18 **IN** `ninfo`
19 Number of elements in `info` (`size_t`)

Description

20 Define a callback function for delivering forwarded IO to a process. This function will be called
21 whenever data becomes available, or a specified buffering size and/or time has been met.
22

Advice to users



23 Multiple strings may be included in a given *payload*, and the *payload* may *not* be **NULL** terminated.
24 The user is responsible for releasing the *payload* memory. The *info* array is owned by the PMIx
25 library and is not to be released or altered by the receiving party.

26 3.5.18 IOF and Event registration function

27 Summary

28 Callback function for calls to register handlers, e.g., event notification and IOF requests.



1 **Format**

2 *PMIx v3.0*  

3 `typedef void (*pmix_hdlr_reg_cbfunc_t) (pmix_status_t status,`

4 `size_t refid,`

5 `void *cbdata);`

6  

5 **IN status**

6 `PMIX_SUCCESS` or an appropriate error constant (`pmix_status_t`)

7 **IN refid**

8 reference identifier assigned to the handler by PMIx, used to deregister the handler (`size_t`)

9 **IN cbdata**

10 object provided to the registration call (pointer)

11 **Description**

12 Callback function for calls to register handlers, e.g., event notification and IOF requests.



13 3.6 Constant String Functions

14 Provide a string representation for several types of values. Note that the provided string is statically

15 defined and must NOT be `free`'d.



16 **Summary**

17 String representation of a `pmix_status_t`.

18 *PMIx v1.0*  



18 `const char*`

19 `PMIx_Error_string (pmix_status_t status);`

20  



20 **Summary**

21 String representation of a `pmix_proc_state_t`.

22 *PMIx v2.0*  

22 `const char*`

23 `PMIx_Proc_state_string (pmix_proc_state_t state);`

24  

1 **Summary**
2 String representation of a `pmix_scope_t`.

PMIx v2.0 ▼  ▼

3 `const char*`
4 `PMIx_Scope_string(pmix_scope_t scope);`

▲  ▲

5 **Summary**
6 String representation of a `pmix_persistence_t`.

PMIx v2.0 ▼  ▼

7 `const char*`
8 `PMIx_Persistence_string(pmix_persistence_t persist);`

▲  ▲

9 **Summary**
10 String representation of a `pmix_data_range_t`.

PMIx v2.0 ▼  ▼

11 `const char*`
12 `PMIx_Data_range_string(pmix_data_range_t range);`

▲  ▲

13 **Summary**
14 String representation of a `pmix_info_directives_t`.

PMIx v2.0 ▼  ▼

15 `const char*`
16 `PMIx_Info_directives_string(pmix_info_directives_t directives);`

▲  ▲

17 **Summary**
18 String representation of a `pmix_data_type_t`.



PMIx v2.0 ▼  ▼

19 `const char*`
20 `PMIx_Data_type_string(pmix_data_type_t type);`

▲  ▲



1 **Summary**
2 String representation of a `pmix_alloc_directive_t`.

PMIx v2.0 ▼  C 

3 **const char***
4 **PMIx_Alloc_directive_string**(`pmix_alloc_directive_t directive`);
 ▲  C 

5 **Summary**
6 String representation of a `pmix_iof_channel_t`.

PMIx v3.0 ▼  C 

7 **const char***
8 **PMIx_IOF_channel_string**(`pmix_iof_channel_t channel`);
 ▲  C 

CHAPTER 4

Initialization and Finalization

1 The PMIx library is required to be initialized and finalized around the usage of most of the APIs.
2 The APIs that may be used outside of the initialized and finalized region are noted. All other APIs
3 must be used inside this region.

4 There are three sets of initialization and finalization functions depending upon the role of the
5 process in the PMIx universe. Each of these functional sets are described in this chapter. Note that
6 a process can only call *one* of the init/finalize functional pairs - e.g., a process that calls the client
7 initialization function cannot also call the tool or server initialization functions, and must call the
8 corresponding client finalize.

Advice to users

9 Processes that initialize as a server or tool automatically are given access to all client APIs. Server
10 initialization includes setting up the infrastructure to support local clients - thus, it necessarily
11 includes overhead and an increased memory footprint. Tool initialization automatically searches for
12 a server to which it can connect — if declared as a *launcher*, the PMIx library sets up the required
13 “hooks” for other tools (e.g., debuggers) to attach to it.

4.1 Query

15 The API defined in this section can be used by any PMIx process, regardless of their role in the
16 PMIx universe.

4.1.1 PMIx_Initialized

Format

PMIx v1.0

```
int PMIx_Initialized(void)
```

20 A value of **1** (true) will be returned if the PMIx library has been initialized, and **0** (false) otherwise.

Rationale

21 The return value is an integer for historical reasons as that was the signature of prior PMI libraries.

1 **Description**
2 Check to see if the PMIx library has been initialized using any of the init functions: `PMIx_Init` ,
3 `PMIx_server_init` , or `PMIx_tool_init` .

4 **4.1.2 PMIx_Get_version**

5 **Summary**
6 Get the PMIx version information.

7 **Format**
8 *PMIx v1.0* ▼ `const char* PMIx_Get_version(void)` ▲ C

9 **Description**
10 Get the PMIx version string. Note that the provided string is statically defined and must *not* be
11 free'd.

12 **4.2 Client Initialization and Finalization**

13 Initialization and finalization routines for PMIx clients.
14 ▼ **Advice to users** ▲
15 The PMIx *ad hoc* v1.0 Standard defined the `PMIx_Init` function, but modified the function
16 signature in the v1.2 version. The *ad hoc* v1.0 version is not included in this document to avoid
confusion.

17 **4.2.1 PMIx_Init**

18 **Summary**
19 Initialize the PMIx client library

1
PMIx v1.2

Format

```

C
pmix_status_t
PMIx_Init (pmix_proc_t *proc,
           pmix_info_t info[], size_t ninfo)
C

```

- 5 **INOUT** `proc`
proc structure (handle)
- 6
- 7 **IN** `info`
Array of `pmix_info_t` structures (array of handles)
- 8
- 9 **IN** `ninfo`
Number of element in the `info` array (`size_t`)
- 10

11 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Optional Attributes

12 The following attributes are optional for implementers of PMIx libraries:

- 13 **PMIX_USOCK_DISABLE** "`pmix.usock.disable`" (`bool`)
14 Disable legacy UNIX socket (usock) support. If the library supports Unix socket
15 connections, this attribute may be supported for disabling it.
- 16 **PMIX_SOCKET_MODE** "`pmix.sockmode`" (`uint32_t`)
17 POSIX `mode_t` (9 bits valid). If the library supports socket connections, this attribute may
18 be supported for setting the socket mode.
- 19 **PMIX_SINGLE_LISTENER** "`pmix.sing.listnr`" (`bool`)
20 Use only one rendezvous socket, letting priorities and/or environment parameters select the
21 active transport. If the library supports multiple methods for clients to connect to servers,
22 this attribute may be supported for disabling all but one of them.
- 23 **PMIX_TCP_REPORT_URI** "`pmix.tcp.repuri`" (`char*`)
24 If provided, directs that the TCP URI be reported and indicates the desired method of
25 reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket
26 connections, this attribute may be supported for reporting the URI.
- 27 **PMIX_TCP_IF_INCLUDE** "`pmix.tcp.ifinclude`" (`char*`)
28 Comma-delimited list of devices and/or CIDR notation to include when establishing the
29 TCP connection. If the library supports TCP socket connections, this attribute may be
30 supported for specifying the interfaces to be used.
- 31 **PMIX_TCP_IF_EXCLUDE** "`pmix.tcp.ifexclude`" (`char*`)
32 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
33 TCP connection. If the library supports TCP socket connections, this attribute may be
34 supported for specifying the interfaces that are *not* to be used.

1 **PMIX_TCP_IPV4_PORT** "pmix.tcp.ipv4" (int)
 2 The IPv4 port to be used. If the library supports IPV4 connections, this attribute may be
 3 supported for specifying the port to be used.

4 **PMIX_TCP_IPV6_PORT** "pmix.tcp.ipv6" (int)
 5 The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be
 6 supported for specifying the port to be used.

7 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)
 8 Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections,
 9 this attribute may be supported for disabling it.

10 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)
 11 Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections,
 12 this attribute may be supported for disabling it.

13 **PMIX_EVENT_BASE** "pmix.evbase" (struct event_base *)
 14 Pointer to libevent¹ **event_base** to use in place of the internal progress thread.

15 **PMIX_GDS_MODULE** "pmix.gds.mod" (char*)
 16 Comma-delimited string of desired modules. This attribute is specific to the PRI and
 17 controls only the selection of GDS module for internal use by the process. Module selection
 18 for interacting with the server is performed dynamically during the connection process.



19 Description

20 Initialize the PMIx client, returning the process identifier assigned to this client's application in the
 21 provided **pmix_proc_t** struct. Passing a value of **NULL** for this parameter is allowed if the user
 22 wishes solely to initialize the PMIx system and does not require return of the identifier at that time.

23 When called, the PMIx client shall check for the required connection information of the local PMIx
 24 server and establish the connection. If the information is not found, or the server connection fails,
 25 then an appropriate error constant shall be returned.

26 If successful, the function shall return **PMIX_SUCCESS** and fill the *proc* structure (if provided)
 27 with the server-assigned namespace and rank of the process within the application. In addition, all
 28 startup information provided by the resource manager shall be made available to the client process
 29 via subsequent calls to **PMIx_Get** .

30 The PMIx client library shall be reference counted, and so multiple calls to **PMIx_Init** are
 31 allowed by the standard. Thus, one way for an application process to obtain its namespace and rank
 32 is to simply call **PMIx_Init** with a non-NULL *proc* parameter. Note that each call to
 33 **PMIx_Init** must be balanced with a call to **PMIx_Finalize** to maintain the reference count.

34 Each call to **PMIx_Init** may contain an array of **pmix_info_t** structures passing directives to
 35 the PMIx client library as per the above attributes.

¹<http://libevent.org/>

1 Multiple calls to `PMIx_Init` shall not include conflicting directives. The `PMIx_Init` function
2 will return an error when directives that conflict with prior directives are encountered.

3 4.2.2 `PMIx_Finalize`

4 Summary

5 Finalize the PMIx client library.

6 Format

PMIx v1.0

```
▼ _____ C _____ ▼  
7 pmix_status_t  
8 PMIx_Finalize(const pmix_info_t info[], size_t ninfo)  
▲ _____ C _____ ▲
```

9 **IN** `info`

10 Array of `pmix_info_t` structures (array of handles)

11 **IN** `ninfo`

12 Number of element in the `info` array (`size_t`)

13 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

▼ ----- Optional Attributes ----- ▼

14 The following attributes are optional for implementers of PMIx libraries:

15 **PMIX_EMBED_BARRIER** "`pmix.embed.barrier`" (`bool`)

16 Execute a blocking fence operation before executing the specified operation. For example,
17 `PMIx_Finalize` does not include an internal barrier operation by default. This attribute
18 would direct `PMIx_Finalize` to execute a barrier as part of the finalize operation.

▲ ----- ▲

19 Description

20 Decrement the PMIx client library reference count. When the reference count reaches zero, the
21 library will finalize the PMIx client, closing the connection with the local PMIx server and
22 releasing all internally allocated memory.

23 4.3 Tool Initialization and Finalization

24 Initialization and finalization routines for PMIx tools.

25 4.3.1 `PMIx_tool_init`

26 Summary

27 Initialize the PMIx library for operating as a tool.

1
PMIx v2.0

Format

```

C
-----▼
pmix_status_t
PMIx_tool_init(pmix_proc_t *proc,
               pmix_info_t info[], size_t ninfo)
-----▲
C

```

INOUT proc

pmix_proc_t structure (handle)

IN info

Array of pmix_info_t structures (array of handles)

IN ninfo

Number of element in the info array (size_t)

Returns PMIX_SUCCESS or a negative value corresponding to a PMIx error constant.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_TOOL_NAMESPACE "pmix.tool.namespace" (char*)

Name of the namespace to use for this tool.

PMIX_TOOL_RANK "pmix.tool.rank" (uint32_t)

Rank of this tool.

PMIX_TOOL_DO_NOT_CONNECT "pmix.tool.nocon" (bool)

The tool wants to use internal PMIx support, but does not want to connect to a PMIx server.

PMIX_SERVER_URI "pmix.srvr.uri" (char*)

URI of the PMIx server to be contacted.

Optional Attributes

The following attributes are optional for implementers of PMIx libraries:

PMIX_CONNECT_TO_SYSTEM "pmix.cnct.sys" (bool)

The requestor requires that a connection be made only to a local, system-level PMIx server.

PMIX_CONNECT_SYSTEM_FIRST "pmix.cnct.sys.first" (bool)

Preferentially, look for a system-level PMIx server first.

PMIX_SERVER_PIDINFO "pmix.srvr.pidinfo" (pid_t)

PID of the target PMIx server for a tool.

PMIX_TCP_URI "pmix.tcp.uri" (char*)

The URI of the PMIx server to connect to, or a file name containing it in the form of file:<name of file containing it>.

PMIX_CONNECT_RETRY_DELAY "pmix.tool.retry" (uint32_t)

1 Time in seconds between connection attempts to a PMI_x server.

2 **PMIX_CONNECT_MAX_RETRIES** "pmix.tool.mretries" (uint32_t)

3 Maximum number of times to try to connect to PMI_x server.

4 **PMIX_SOCKET_MODE** "pmix.sockmode" (uint32_t)

5 POSIX *mode_t* (9 bits valid) If the library supports socket connections, this attribute may
6 be supported for setting the socket mode.

7 **PMIX_TCP_REPORT_URI** "pmix.tcp.repuri" (char*)

8 If provided, directs that the TCP URI be reported and indicates the desired method of
9 reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket
10 connections, this attribute may be supported for reporting the URI.

11 **PMIX_TCP_IF_INCLUDE** "pmix.tcp.ifinclude" (char*)

12 Comma-delimited list of devices and/or CIDR notation to include when establishing the
13 TCP connection. If the library supports TCP socket connections, this attribute may be
14 supported for specifying the interfaces to be used.

15 **PMIX_TCP_IF_EXCLUDE** "pmix.tcp.ifexclude" (char*)

16 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
17 TCP connection. If the library supports TCP socket connections, this attribute may be
18 supported for specifying the interfaces that are *not* to be used.

19 **PMIX_TCP_IPV4_PORT** "pmix.tcp.ipv4" (int)

20 The IPv4 port to be used. If the library supports IPV4 connections, this attribute may be
21 supported for specifying the port to be used.

22 **PMIX_TCP_IPV6_PORT** "pmix.tcp.ipv6" (int)

23 The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be
24 supported for specifying the port to be used.

25 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)

26 Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections,
27 this attribute may be supported for disabling it.

28 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)

29 Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections,
30 this attribute may be supported for disabling it.

31 **PMIX_EVENT_BASE** "pmix.evbase" (struct event_base *)

32 Pointer to libevent² **event_base** to use in place of the internal progress thread.

33 **PMIX_GDS_MODULE** "pmix.gds.mod" (char*)

34 Comma-delimited string of desired modules. This attribute is specific to the PRI and
35 controls only the selection of GDS module for internal use by the process. Module selection
36 for interacting with the server is performed dynamically during the connection process.

²<http://libevent.org/>



Description

1 Initialize the PMIx tool, returning the process identifier assigned to this tool in the provided
2 `pmix_proc_t` struct. The *info* array is used to pass user requests pertaining to the init and
3 subsequent operations. Passing a **NULL** value for the array pointer is supported if no directives are
4 desired.
5

6 If called with the **PMIX_TOOL_DO_NOT_CONNECT** attribute, the PMIx tool library will fully
7 initialize but not attempt to connect to a PMIx server. The tool can connect to a server at a later
8 point in time, if desired. In all other cases, the PMIx tool library will attempt to connect to
9 according to the following precedence chain:

- 10 • if **PMIX_SERVER_URI** or **PMIX_TCP_URI** is given, then connection will be attempted to the
11 server at the specified URI. Note that it is an error for both of these attributes to be specified.
12 **PMIX_SERVER_URI** is the preferred method as it is more generalized — **PMIX_TCP_URI** is
13 provided for those cases where the user specifically wants to use a TCP transport for the
14 connection and wants to error out if it isn't available or cannot succeed. The PMIx library will
15 return an error if connection fails — it will not proceed to check for other connection options as
16 the user specified a particular one to use
- 17 • if **PMIX_SERVER_PIDINFO** was provided, then the tool will search under the directory
18 provided by the `PMIX_SERVER_TMPDIR` environmental variable for a rendezvous file created
19 by the process corresponding to that PID. The PMIx library will return an error if the rendezvous
20 file cannot be found, or the connection is refused by the server
- 21 • if **PMIX_CONNECT_TO_SYSTEM** is given, then the tool will search for a system-level
22 rendezvous file created by a PMIx server in the directory specified by the
23 `PMIX_SYSTEM_TMPDIR` environmental variable. If found, then the tool will attempt to
24 connect to it. An error is returned if the rendezvous file cannot be found or the connection is
25 refused.
- 26 • if **PMIX_CONNECT_SYSTEM_FIRST** is given, then the tool will search for a system-level
27 rendezvous file created by a PMIx server in the directory specified by the
28 `PMIX_SYSTEM_TMPDIR` environmental variable. If found, then the tool will attempt to
29 connect to it. In this case, no error will be returned if the rendezvous file is not found or
30 connection is refused — the PMIx library will silently continue to the next option
- 31 • by default, the tool will search the directory tree under the directory provided by the
32 `PMIX_SERVER_TMPDIR` environmental variable for rendezvous files of PMIx servers,
33 attempting to connect to each it finds until one accepts the connection. If no rendezvous files are
34 found, or all contacted servers refuse connection, then the PMIx library will return an error.

35 If successful, the function will return **PMIX_SUCCESS** and will fill the provided structure (if
36 provided) with the server-assigned namespace and rank of the tool. Note that each connection
37 attempt in the above precedence chain will retry (with delay between each retry) a number of times
38 according to the values of the corresponding attributes. Default is no retries.

1 Note that the PMIx tool library is referenced counted, and so multiple calls to `PMIx_tool_init`
2 are allowed. Thus, one way to obtain the namespace and rank of the process is to simply call
3 `PMIx_tool_init` with a non-NULL parameter.

4 4.3.2 `PMIx_tool_finalize`

5 **Summary**

6 Finalize the PMIx library for a tool connection.

7 **Format**

PMIx v2.0

```
▼ _____ C _____ ▼  
pmix_status_t  
PMIx_tool_finalize(void)  
▲ _____ C _____ ▲
```

10 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

11 **Description**

12 Finalize the PMIx tool library, closing the connection to the server. An error code will be returned
13 if, for some reason, the connection cannot be cleanly terminated — in this case, the connection is
14 dropped.

15 4.3.3 `PMIx_tool_connect_to_server`

16 **Summary**

17 Switch connection from the current PMIx server to another one, or initialize a connection to a
18 specified server.

19 **Format**

PMIx v3.0

```
▼ _____ C _____ ▼  
pmix_status_t  
PMIx_tool_connect_to_server(pmix_proc_t *proc,  
                             pmix_info_t info[], size_t ninfo)  
▲ _____ C _____ ▲
```

23 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_CONNECT_TO_SYSTEM "pmix.cnct.sys" (bool)

The requestor requires that a connection be made only to a local, system-level PMIx server.

PMIX_CONNECT_SYSTEM_FIRST "pmix.cnct.sys.first" (bool)

Preferentially, look for a system-level PMIx server first.

PMIX_SERVER_URI "pmix.srvr.uri" (char*)

URI of the PMIx server to be contacted.

PMIX_SERVER_NAMESPACE "pmix.srv.namespace" (char*)

Name of the namespace to use for this PMIx server.

PMIX_SERVER_PIDINFO "pmix.srvr.pidinfo" (pid_t)

PID of the target PMIx server for a tool.

Description

Switch connection from the current PMIx server to another one, or initialize a connection to a specified server. Closes the connection, if existing, to a server and establishes a connection to the specified server. This function can be called at any time by a PMIx tool to shift connections between servers. The process identifier assigned to this tool is returned in the provided `pmix_proc_t` struct. Passing a value of **NULL** for this parameter is allowed if the user wishes solely to connect to the PMIx server and does not require return of the identifier at that time.

Advice to PMIx library implementers

PMIx tools and clients are prohibited from being connected to more than one server at a time to avoid confusion in subsystems such as event notification.

When a tool connects to a server that is under a different namespace manager (e.g., host RM) as the prior server, the identifier of the tool must remain unique in the namespaces. This may require the identifier of the tool to be changed on-the-fly, that is, the *proc* parameter would be filled (if non-NULL) with a different namespace/rank from the current tool identifier.

Advice to users

Passing a **NULL** value for the *info* pointer is not allowed and will result in returning an error.

Some PMIx implementations (for example, the current PRI) may not support connecting to a server that is not under the same namespace manager (e.g., host RM) as the tool.

1 4.4 Server Initialization and Finalization

2 Initialization and finalization routines for PMIx servers.

3 4.4.1 PMIx_server_init

4 Summary

5 Initialize the PMIx server.

6 Format

PMIx v1.0

```
▼----- C -----▼  
pmix_status_t  
PMIx_server_init(pmix_server_module_t *module,  
                 pmix_info_t info[], size_t ninfo)  
▲----- C -----▲
```

10 INOUT module

11 `pmix_server_module_t` structure (handle)

12 IN info

13 Array of `pmix_info_t` structures (array of handles)

14 IN ninfo

15 Number of elements in the *info* array (`size_t`)

16 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

▼----- Required Attributes -----▼

17 The following attributes are required to be supported by all PMIx libraries:

18 **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (`char*`)

19 Name of the namespace to use for this PMIx server.

20 **PMIX_SERVER_RANK** "pmix.srv.rank" (`pmix_rank_t`)

21 Rank of this PMIx server

22 **PMIX_SERVER_TMPDIR** "pmix.srvr.tmpdir" (`char*`)

23 Top-level temporary directory for all client processes connected to this server, and where the
24 PMIx server will place its tool rendezvous point and contact information.

25 **PMIX_SYSTEM_TMPDIR** "pmix.sys.tmpdir" (`char*`)

26 Temporary directory for this system, and where a PMIx server that declares itself to be a
27 system-level server will place a tool rendezvous point and contact information.

28 **PMIX_SERVER_TOOL_SUPPORT** "pmix.srvr.tool" (`bool`)

29 The host RM wants to declare itself as willing to accept tool connection requests.

30 **PMIX_SERVER_SYSTEM_SUPPORT** "pmix.srvr.sys" (`bool`)

31 The host RM wants to declare itself as being the local system server for PMIx connection
32 requests.

Optional Attributes

The following attributes are optional for implementers of PMIx libraries:

PMIX_USOCK_DISABLE "pmix.usock.disable" (bool)

Disable legacy UNIX socket (usock) support. If the library supports Unix socket connections, this attribute may be supported for disabling it.

PMIX_SOCKET_MODE "pmix.sockmode" (uint32_t)

POSIX *mode_t* (9 bits valid). If the library supports socket connections, this attribute may be supported for setting the socket mode.

PMIX_TCP_REPORT_URI "pmix.tcp.repuri" (char*)

If provided, directs that the TCP URI be reported and indicates the desired method of reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket connections, this attribute may be supported for reporting the URI.

PMIX_TCP_IF_INCLUDE "pmix.tcp.ifinclude" (char*)

Comma-delimited list of devices and/or CIDR notation to include when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces to be used.

PMIX_TCP_IF_EXCLUDE "pmix.tcp.ifexclude" (char*)

Comma-delimited list of devices and/or CIDR notation to exclude when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces that are *not* to be used.

PMIX_TCP_IPV4_PORT "pmix.tcp.ipv4" (int)

The IPv4 port to be used. If the library supports IPV4 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_IPV6_PORT "pmix.tcp.ipv6" (int)

The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_DISABLE_IPV4 "pmix.tcp.disipv4" (bool)

Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections, this attribute may be supported for disabling it.

PMIX_TCP_DISABLE_IPV6 "pmix.tcp.disipv6" (bool)

Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections, this attribute may be supported for disabling it.

PMIX_SERVER_REMOTE_CONNECTIONS "pmix.srvr.remote" (bool)

Allow connections from remote tools. Forces the PMIx server to not exclusively use loopback device. If the library supports connections from remote tools, this attribute may be supported for enabling or disabling it.

1 **PMIX_EVENT_BASE** "pmix.evbase" (**struct event_base ***)
 2 Pointer to libevent³ **event_base** to use in place of the internal progress thread.

3 **PMIX_GDS_MODULE** "pmix.gds.mod" (**char***)
 4 Comma-delimited string of desired modules. This attribute is specific to the PRI and
 5 controls only the selection of GDS module for internal use by the process. Module selection
 6 for interacting with the server is performed dynamically during the connection process.



7 **Description**
 8 Initialize the PMIx server support library, and provide a pointer to a **pmix_server_module_t**
 9 structure containing the caller's callback functions. The array of **pmix_info_t** structs is used to
 10 pass additional info that may be required by the server when initializing. For example, it may
 11 include the **PMIX_SERVER_TOOL_SUPPORT** attribute, thereby indicating that the daemon is
 12 willing to accept connection requests from tools.

▼ **Advice to PMIx server hosts** ▼

13 Providing a value of **NULL** for the *module* argument is permitted, as is passing an empty *module*
 14 structure. Doing so indicates that the host environment will not provide support for multi-node
 15 operations such as **PMIx_Fence**, but does intend to support local clients access to information.



16 **4.4.2 PMIx_server_finalize**

17 **Summary**
 18 Finalize the PMIx server library.

19 **Format**
 20 *PMIx v1.0* ▼ **C** ▼
 21 **pmix_status_t**
PMIx_server_finalize(void) **C** ▲

22 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

23 **Description**
 24 Finalize the PMIx server support library, terminating all connections to attached tools and any local
 25 clients. All memory usage is released.

³<http://libevent.org/>

CHAPTER 5

Key/Value Management

1 Management of key-value pairs in PMIx is a distributed responsibility. While the stated objective of
2 the PMIx community is to eliminate collective operations, it is recognized that the traditional
3 method of posting/exchanging data must be supported until that objective can be met. This method
4 relies on processes to discover and post their local information which is collected by the local PMIx
5 server library. Global exchange of the posted information is then executed via a collective operation
6 performed by the host SMS servers. The `PMIx_Put` and `PMIx_Commit` APIs, plus an attribute
7 directing `PMIx_Fence` to globally collect the data posted by processes, are provided for this
8 purpose.

5.1 Setting and Accessing Key/Value Pairs

5.1.1 `PMIx_Put`

Summary

Push a key/value pair into the client's namespace.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_Put (pmix_scope_t scope,  
          const pmix_key_t key,  
          pmix_value_t *val)
```

IN scope

Distribution scope of the provided value (handle)

IN key

key (`pmix_key_t`)

IN value

Reference to a `pmix_value_t` structure (handle)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Push a value into the client's namespace. The client's PMIx library will cache the information locally until `PMIx_Commit` is called.

The provided *scope* is passed to the local PMIx server, which will distribute the data to other processes according to the provided scope. The `pmix_scope_t` values are defined in Section 3.2.9 on page 33. Specific implementations may support different scope values, but all implementations must support at least `PMIX_GLOBAL`.

The `pmix_value_t` structure supports both string and binary values. PMIx implementations will support heterogeneous environments by properly converting binary values between host architectures, and will copy the provided *value* into internal memory.

Advice to PMIx library implementers

The PMIx server library will properly pack/unpack data to accommodate heterogeneous environments. The host SMS is not involved in this action. The *value* argument must be copied - the caller is free to release it following return from the function.

Advice to users

The value is copied by the PMIx client library. Thus, the application is free to release and/or modify the value once the call to `PMIx_Put` has completed.

Note that keys starting with a string of "`pmix`" are exclusively reserved for the PMIx standard and must not be used in calls to `PMIx_Put`. Thus, applications should never use a defined "`PMIX_`" attribute as the key in a call to `PMIx_Put`.

5.1.2 `PMIx_Get`

Summary

Retrieve a key/value pair from the client's namespace.

1
PMIx v1.0

Format

C

```
2 pmix_status_t  
3 PMIx_Get(const pmix_proc_t *proc, const pmix_key_t key,  
4         const pmix_info_t info[], size_t ninfo,  
5         pmix_value_t **val)
```

C

6 **IN** **proc**
7 process reference (handle)
8 **IN** **key**
9 key to retrieve ([pmix_key_t](#))
10 **IN** **info**
11 Array of info structures (array of handles)
12 **IN** **ninfo**
13 Number of element in the *info* array (integer)
14 **OUT** **val**
15 value (handle)

16 Returns [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant.

Required Attributes

17 The following attributes are required to be supported by all PMIx libraries:

- 18 **PMIX_OPTIONAL** "pmix.optional" (bool)
19 Look only in the client's local data store for the requested value - do not request data from
20 the PMIx server if not found.
- 21 **PMIX_IMMEDIATE** "pmix.immediate" (bool)
22 Specified operation should immediately return an error from the PMIx server if the requested
23 data cannot be found - do not request it from the host RM.
- 24 **PMIX_DATA_SCOPE** "pmix.scope" ([pmix_scope_t](#))
25 Scope of the data to be found in a [PMIx_Get](#) call.
- 26 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)
27 Return information about the specified session. If information about a session other than the
28 one containing the requesting process is desired, then the attribute array must contain a
29 [PMIX_SESSION_ID](#) attribute identifying the desired target.
- 30 **PMIX_JOB_INFO** "pmix.job.info" (bool)

1 Return information about the specified job or namespace. If information about a job or
2 namespace other than the one containing the requesting process is desired, then the attribute
3 array must contain a `PMIX_JOBID` or `PMIX_NAMESPACE` attribute identifying the desired
4 target. Similarly, if information is requested about a job or namespace in a session other than
5 the one containing the requesting process, then an attribute identifying the target session
6 must be provided.

7 **PMIX_APP_INFO** `"pmix.app.info"` (bool)

8 Return information about the specified application. If information about an application other
9 than the one containing the requesting process is desired, then the attribute array must
10 contain a `PMIX_APPNUM` attribute identifying the desired target. Similarly, if information is
11 requested about an application in a job or session other than the one containing the requesting
12 process, then attributes identifying the target job and/or session must be provided.

13 **PMIX_NODE_INFO** `"pmix.node.info"` (bool)

14 Return information about the specified node. If information about a node other than the one
15 containing the requesting process is desired, then the attribute array must contain either the
16 `PMIX_NODEID` or `PMIX_HOSTNAME` attribute identifying the desired target.



Optional Attributes



17 The following attributes are optional for host environments:

18 **PMIX_TIMEOUT** `"pmix.timeout"` (int)

19 Time in seconds before the specified operation should time out (*0* indicating infinite) in
20 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
21 the target process from ever exposing its data.



Advice to PMIx library implementers



22 We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host
23 environment due to race condition considerations between delivery of the data by the host
24 environment versus internal timeout in the PMIx server library. Implementers that choose to
25 support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race
26 condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple
27 competing timeouts are not created.



Description

Retrieve information for the specified *key* as published by the process identified in the given `pmix_proc_t`, returning a pointer to the value in the given address.

This is a blocking operation - the caller will block until either the specified data becomes available from the specified rank in the *proc* structure or the operation times out should the `PMIX_TIMEOUT` attribute have been given. The caller is responsible for freeing all memory associated with the returned *value* when no longer required.

The *info* array is used to pass user requests regarding the get operation.

Advice to users

Information provided by the PMIx server at time of process start is accessed by providing the namespace of the job with the rank set to `PMIX_RANK_WILDCARD`. The list of data referenced in this way is maintained on the PMIx web site at <https://pmix.org/support/faq/wildcard-rank-access/> but includes items such as the number of processes in the namespace (`PMIX_JOB_SIZE`), total available slots in the allocation (`PMIX_UNIV_SIZE`), and the number of nodes in the allocation (`PMIX_NUM_NODES`).

Data posted by a process via `PMIx_Put` needs to be retrieved by specifying the rank of the posting process. All other information is retrievable using a rank of `PMIX_RANK_WILDCARD` when the information being retrieved refers to something non-rank specific (e.g., number of processes on a node, number of processes in a job), and using the rank of the relevant process when requesting information that is rank-specific (e.g., the URI of the process, or the node upon which it is executing). Each subsection of Section 3.4 indicates the appropriate rank value for referencing the defined attribute.

5.1.3 PMIx_Get_nb

Summary

Nonblocking `PMIx_Get` operation.

1
PMIx v1.0

Format

C

```

2 pmix_status_t
3 PMIx_Get_nb(const pmix_proc_t *proc, const char key[],
4             const pmix_info_t info[], size_t ninfo,
5             pmix_value_cbfunc_t cbfunc, void *cbdata)

```

C

- 6 **IN proc**
process reference (handle)
- 7
- 8 **IN key**
key to retrieve (string)
- 9
- 10 **IN info**
Array of info structures (array of handles)
- 11
- 12 **IN ninfo**
Number of elements in the *info* array (integer)
- 13
- 14 **IN cbfunc**
Callback function (function reference)
- 15
- 16 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 17

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
21 function prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will *not* be called
- 24 • a PMIx error constant indicating either an error in the input or that the request was immediately
25 processed and failed - the *cbfunc* will *not* be called

26 If executed, the status returned in the provided callback function will be one of the following
27 constants:

- 28 • **PMIX_SUCCESS** The requested data has been returned
- 29 • **PMIX_ERR_NOT_FOUND** The requested data was not available
- 30 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

31 The following attributes are required to be supported by all PMIx libraries:

- 32 **PMIX_OPTIONAL** "pmix.optional" (bool)
33 Look only in the client's local data store for the requested value - do not request data from
34 the PMIx server if not found.

1 **PMIX_IMMEDIATE** "pmix.immediate" (bool)
2 Specified operation should immediately return an error from the PMIx server if the requested
3 data cannot be found - do not request it from the host RM.

4 **PMIX_DATA_SCOPE** "pmix.scope" (pmix_scope_t)
5 Scope of the data to be found in a **PMIx_Get** call.

6 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)
7 Return information about the specified session. If information about a session other than the
8 one containing the requesting process is desired, then the attribute array must contain a
9 **PMIX_SESSION_ID** attribute identifying the desired target.

10 **PMIX_JOB_INFO** "pmix.job.info" (bool)
11 Return information about the specified job or namespace. If information about a job or
12 namespace other than the one containing the requesting process is desired, then the attribute
13 array must contain a **PMIX_JOBID** or **PMIX_NAMESPACE** attribute identifying the desired
14 target. Similarly, if information is requested about a job or namespace in a session other than
15 the one containing the requesting process, then an attribute identifying the target session
16 must be provided.

17 **PMIX_APP_INFO** "pmix.app.info" (bool)
18 Return information about the specified application. If information about an application other
19 than the one containing the requesting process is desired, then the attribute array must
20 contain a **PMIX_APPNUM** attribute identifying the desired target. Similarly, if information is
21 requested about an application in a job or session other than the one containing the requesting
22 process, then attributes identifying the target job and/or session must be provided.

23 **PMIX_NODE_INFO** "pmix.node.info" (bool)
24 Return information about the specified node. If information about a node other than the one
25 containing the requesting process is desired, then the attribute array must contain either the
26 **PMIX_NODEID** or **PMIX_HOSTNAME** attribute identifying the desired target.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

27 The following attributes are optional for host environments that support this operation:

28 **PMIX_TIMEOUT** "pmix.timeout" (int)
29 Time in seconds before the specified operation should time out (0 indicating infinite) in
30 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
31 the target process from ever exposing its data.

▲-----▲

Advice to PMIx library implementers

1 We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host
2 environment due to race condition considerations between delivery of the data by the host
3 environment versus internal timeout in the PMIx server library. Implementers that choose to
4 support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race
5 condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple
6 competing timeouts are not created.

Description

7 The callback function will be executed once the specified data becomes available from the
8 identified process and retrieved by the local server. The *info* array is used as described by the
9 `PMIx_Get` routine.
10

Advice to users

11 Information provided by the PMIx server at time of process start is accessed by providing the
12 namespace of the job with the rank set to `PMIX_RANK_WILDCARD`. Attributes referenced in this
13 way are identified in 3.4 but includes items such as the number of processes in the namespace (
14 `PMIX_JOB_SIZE`), total available slots in the allocation (`PMIX_UNIV_SIZE`), and the number
15 of nodes in the allocation (`PMIX_NUM_NODES`).

16 In general, data posted by a process via `PMIx_Put` and data that refers directly to a
17 process-related value needs to be retrieved by specifying the rank of the posting process. All other
18 information is retrievable using a rank of `PMIX_RANK_WILDCARD`, as illustrated in 5.1.5. See
19 3.4.11 for an explanation regarding use of the *level* attributes.

20 5.1.4 PMIx_Store_internal

21 Summary

22 Store some data locally for retrieval by other areas of the proc.

1
PMIx v1.0

Format

C

```
2 pmix_status_t
3 PMIx_Store_internal(const pmix_proc_t *proc,
4                    const pmix_key_t key,
5                    pmix_value_t *val);
```

C

```
6 IN proc
7     process reference (handle)
8 IN key
9     key to retrieve (string)
10 IN val
11     Value to store (handle)
```

12 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Description

14 Store some data locally for retrieval by other areas of the proc. This is data that has only internal
15 scope - it will never be “pushed” externally.

16 5.1.5 Accessing information: examples

17 This section provides examples illustrating methods for accessing information at various levels.
18 The intent of the examples is not to provide comprehensive coding guidance, but rather to illustrate
19 how **PMIx_Get** can be used to obtain information on a **session**, **job**, **application**,
20 process, and node.

21 5.1.5.1 Session-level information

22 The **PMIx_Get** API does not include an argument for specifying the **session** associated with
23 the information being requested. Information regarding the session containing the requestor can be
24 obtained by the following methods:

- 25 • for session-level attributes (e.g., **PMIX_UNIV_SIZE**), specifying the requestor’s namespace
26 and a rank of **PMIX_RANK_WILDCARD**; or
- 27 • for non-specific attributes (e.g., **PMIX_NUM_NODES**), including the **PMIX_SESSION_INFO**
28 attribute to indicate that the session-level information for that attribute is being requested

29 Example requests are shown below:

C

```

1  pmix_info_t info;
2  pmix_value_t *value;
3  pmix_status_t rc;
4  pmix_proc_t myproc, wildcard;
5
6  /* initialize the client library */
7  PMIx_Init(&myproc, NULL, 0);
8
9  /* get the #slots in our session */
10 PMIX_PROC_LOAD(&wildcard, myproc.nspace, PMIX_RANK_WILDCARD);
11 rc = PMIx_Get(&wildcard, PMIX_UNIV_SIZE, NULL, 0, &value);
12
13 /* get the #nodes in our session */
14 PMIX_INFO_LOAD(&info, PMIX_SESSION_INFO, NULL, PMIX_BOOL);
15 rc = PMIx_Get(&wildcard, PMIX_NUM_NODES, &info, 1, &value);

```

C

Information regarding a different session can be requested by either specifying the namespace and a rank of `PMIX_RANK_WILDCARD` for a process in the target session, or adding the `PMIX_SESSION_ID` attribute identifying the target session. In the latter case, the *proc* argument to `PMIx_Get` will be ignored:

C

```

20 pmix_info_t info[2];
21 pmix_value_t *value;
22 pmix_status_t rc;
23 pmix_proc_t myproc;
24 uint32_t sid;
25
26 /* initialize the client library */
27 PMIx_Init(&myproc, NULL, 0);
28
29 /* get the #nodes in a different session */
30 sid = 12345;
31 PMIX_INFO_LOAD(&info[0], PMIX_SESSION_INFO, NULL, PMIX_BOOL);
32 PMIX_INFO_LOAD(&info[1], PMIX_SESSION_ID, &sid, PMIX_UINT32);
33 rc = PMIx_Get(&myproc, PMIX_NUM_NODES, info, 2, &value);

```

C

1 5.1.5.2 Job-level information

2 Information regarding a job can be obtained by the following methods:

- 3 • for job-level attributes (e.g., `PMIX_JOB_SIZE` or `PMIX_JOB_NUM_APPS`), specifying the
4 namespace of the job and a rank of `PMIX_RANK_WILDCARD` for the *proc* argument to
5 `PMIx_Get` ; or
- 6 • for non-specific attributes (e.g., `PMIX_NUM_NODES`), including the `PMIX_JOB_INFO`
7 attribute to indicate that the job-level information for that attribute is being requested

8 Example requests are shown below:

```
9 pmix_info_t info;  
10 pmix_value_t *value;  
11 pmix_status_t rc;  
12 pmix_proc_t myproc, wildcard;  
13  
14 /* initialize the client library */  
15 PMIx_Init(&myproc, NULL, 0);  
16  
17 /* get the #apps in our job */  
18 PMIX_PROC_LOAD(&wildcard, myproc.namespace, PMIX_RANK_WILDCARD);  
19 rc = PMIx_Get(&wildcard, PMIX_JOB_NUM_APPS, NULL, 0, &value);  
20  
21 /* get the #nodes in our job */  
22 PMIX_INFO_LOAD(&info, PMIX_JOB_INFO, NULL, PMIX_BOOL);  
23 rc = PMIx_Get(&wildcard, PMIX_NUM_NODES, &info, 1, &value);
```

24 5.1.5.3 Application-level information

25 Information regarding an application can be obtained by the following methods:

- 26 • for application-level attributes (e.g., `PMIX_APP_SIZE`), specifying the namespace and rank of
27 a process within that application;
- 28 • for application-level attributes (e.g., `PMIX_APP_SIZE`), including the `PMIX_APPNUM`
29 attribute specifying the application whose information is being requested. In this case, the
30 namespace field of the *proc* argument is used to reference the *job* containing the application -
31 the *rank* field is ignored;
- 32 • or application-level attributes (e.g., `PMIX_APP_SIZE`), including the `PMIX_APPNUM` and
33 `PMIX_NAMESPACE` or `PMIX_JOBID` attributes specifying the job/application whose information
34 is being requested. In this case, the *proc* argument is ignored;
- 35 • for non-specific attributes (e.g., `PMIX_NUM_NODES`), including the `PMIX_APP_INFO`
36 attribute to indicate that the application-level information for that attribute is being requested

1 Example requests are shown below:

```
2 pmix_info_t info;
3 pmix_value_t *value;
4 pmix_status_t rc;
5 pmix_proc_t myproc, otherproc;
6 uint32_t appsize, appnum;
7
8 /* initialize the client library */
9 PMIx_Init(&myproc, NULL, 0);
10
11 /* get the #processes in our application */
12 rc = PMIx_Get(&myproc, PMIX_APP_SIZE, NULL, 0, &value);
13 appsize = value->data.uint32;
14
15 /* get the #nodes in an application containing "otherproc".
16  * Note that the rank of a process in the other application
17  * must be obtained first - a simple method is shown here */
18
19 /* assume for this example that we are in the first application
20  * and we want the #nodes in the second application - use the
21  * rank of the first process in that application, remembering
22  * that ranks start at zero */
23 PMIX_PROC_LOAD(&otherproc, myproc.namespace, appsize);
24
25 PMIX_INFO_LOAD(&info, PMIX_APP_INFO, NULL, PMIX_BOOL);
26 rc = PMIx_Get(&otherproc, PMIX_NUM_NODES, &info, 1, &value);
27
28 /* alternatively, we can directly ask for the #nodes in
29  * the second application in our job, again remembering that
30  * application numbers start with zero */
31 appnum = 1;
32 PMIX_INFO_LOAD(&appinfo[0], PMIX_APP_INFO, NULL, PMIX_BOOL);
33 PMIX_INFO_LOAD(&appinfo[1], PMIX_APPNUM, &appnum, PMIX_UINT32);
34 rc = PMIx_Get(&myproc, PMIX_NUM_NODES, appinfo, 2, &value);
35
```

36 5.1.5.4 Process-level information

37 Process-level information is accessed by providing the namespace and rank of the target process. In
38 the absence of any directive as to the level of information being requested, the PMIx library will
39 always return the process-level value.

1 5.1.5.5 Node-level information

2 Information regarding a node within the system can be obtained by the following methods:

- 3 • for node-level attributes (e.g., [PMIX_NODE_SIZE](#)), specifying the namespace and rank of a
4 process executing on the target node;
- 5 • for node-level attributes (e.g., [PMIX_NODE_SIZE](#)), including the [PMIX_NODEID](#) or
6 [PMIX_HOSTNAME](#) attribute specifying the node whose information is being requested. In this
7 case, the *proc* argument's values are ignored; or
- 8 • for non-specific attributes (e.g., [PMIX_MAX_PROCS](#)), including the [PMIX_NODE_INFO](#)
9 attribute to indicate that the node-level information for that attribute is being requested

10 Example requests are shown below:

```
11 pmix_info_t info[2];  
12 pmix_value_t *value;  
13 pmix_status_t rc;  
14 pmix_proc_t myproc, otherproc;  
15 uint32_t nodeid;  
  
16 /* initialize the client library */  
17 PMIx_Init(&myproc, NULL, 0);  
18  
19  
20 /* get the #procs on our node */  
21 rc = PMIx_Get(&myproc, PMIX_NODE_SIZE, NULL, 0, &value);  
22  
23 /* get the #slots on another node */  
24 PMIX_INFO_LOAD(&info[0], PMIX_NODE_INFO, NULL, PMIX_BOOL);  
25 PMIX_INFO_LOAD(&info[1], PMIX_HOSTNAME, "remotehost", PMIX_STRING);  
26 rc = PMIx_Get(&myproc, PMIX_MAX_PROCS, info, 2, &value);  
27
```

28 An explanation of the use of [PMIx_Get](#) versus [PMIx_Query_info_nb](#) is provided in [7.1.4.1](#).

29 5.2 Exchanging Key/Value Pairs

30 The APIs defined in this section push key/value pairs from the client to the local PMIx server, and
31 circulate the data between PMIx servers for subsequent retrieval by the local clients.

1 5.2.1 PMIx_Commit

2 Summary

3 Push all previously `PMIx_Put` values to the local PMIx server.

4 Format

PMIx v1.0

C

5 `pmix_status_t PMIx_Commit(void)`

C

6 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

7 Description

8 This is an asynchronous operation. The PRI will immediately return to the caller while the data is
9 transmitted to the local server in the background.

Advice to users

10 The local PMIx server will cache the information locally - i.e., the committed data will not be
11 circulated during `PMIx_Commit`. Availability of the data upon completion of `PMIx_Commit` is
12 therefore implementation-dependent.

13 5.2.2 PMIx_Fence

14 Summary

15 Execute a blocking barrier across the processes identified in the specified array, collecting
16 information posted via `PMIx_Put` as directed.

1
PMIx v1.0

Format

C

```
2 pmix_status_t
3 PMIx_Fence(const pmix_proc_t procs[], size_t nprocs,
4             const pmix_info_t info[], size_t ninfo)
```

C

- 5 **IN procs**
Array of `pmix_proc_t` structures (array of handles)
- 7 **IN nprocs**
Number of element in the `procs` array (integer)
- 9 **IN info**
Array of info structures (array of handles)
- 11 **IN ninfo**
Number of element in the `info` array (integer)

13 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Required Attributes

14 The following attributes are required to be supported by all PMIx libraries:

15 **PMIX_COLLECT_DATA** "pmix.collect" (bool)
16 Collect data and return it at the end of the operation.

Optional Attributes

17 The following attributes are optional for host environments:

18 **PMIX_TIMEOUT** "pmix.timeout" (int)
19 Time in seconds before the specified operation should time out (0 indicating infinite) in
20 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
21 the target process from ever exposing its data.

22 **PMIX_COLLECTIVE_ALGO** "pmix.calgo" (char*)
23 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
24 impose any requirements on a host environment's collective algorithms. Thus, the
25 acceptable values for this attribute will be environment-dependent - users are encouraged to
26 check their host environment for supported values.

27 **PMIX_COLLECTIVE_ALGO_REQD** "pmix.calreqd" (bool)
28 If `true`, indicates that the requested choice of algorithm is mandatory.

Advice to PMIx library implementers

1 We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host
2 environment due to race condition considerations between completion of the operation versus
3 internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT`
4 directly in the PMIx server library must take care to resolve the race condition and should avoid
5 passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not
6 created.

Description

7 Passing a `NULL` pointer as the *procs* parameter indicates that the fence is to span all processes in
8 the client's namespace. Each provided `pmix_proc_t` struct can pass `PMIX_RANK_WILDCARD`
9 to indicate that all processes in the given namespace are participating.

11 The *info* array is used to pass user requests regarding the fence operation.

12 Note that for scalability reasons, the default behavior for `PMIx_Fence` is to not collect the data.

Advice to PMIx library implementers

13 `PMIx_Fence` and its non-blocking form are both *collective* operations. Accordingly, the PMIx
14 server library is required to aggregate participation by local clients, passing the request to the host
15 environment once all local participants have executed the API.

Advice to PMIx server hosts

16 The host will receive a single call for each collective operation. It is the responsibility of the host to
17 identify the nodes containing participating processes, execute the collective across all participating
18 nodes, and notify the local PMIx server library upon completion of the global collective.

5.2.3 `PMIx_Fence_nb`

Summary

20 Execute a nonblocking `PMIx_Fence` across the processes identified in the specified array of
21 processes, collecting information posted via `PMIx_Put` as directed.
22

1
PMIx v1.0

Format

C

```

2 pmix_status_t
3 PMIx_Fence_nb(const pmix_proc_t procs[], size_t nprocs,
4               const pmix_info_t info[], size_t ninfo,
5               pmix_op_cbfunc_t cbfunc, void *cbdata)

```

C

- 6 **IN procs**
Array of `pmix_proc_t` structures (array of handles)
- 7
- 8 **IN nprocs**
Number of element in the *procs* array (integer)
- 9
- 10 **IN info**
Array of info structures (array of handles)
- 11
- 12 **IN ninfo**
Number of element in the *info* array (integer)
- 13
- 14 **IN cbfunc**
Callback function (function reference)
- 15
- 16 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 17

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
21 function prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will *not* be called. This can occur if the collective involved only
24 processes on the local node.
- 25 • a PMIx error constant indicating either an error in the input or that the request was immediately
26 processed and failed - the *cbfunc* will *not* be called

Required Attributes

27 The following attributes are required to be supported by all PMIx libraries:

- 28 **PMIX_COLLECT_DATA** "pmix.collect" (bool)
29 Collect data and return it at the end of the operation.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_COLLECTIVE_ALGO "pmix.calgo" (char*)

Comma-delimited list of algorithms to use for the collective operation. PMIx does not impose any requirements on a host environment’s collective algorithms. Thus, the acceptable values for this attribute will be environment-dependent - users are encouraged to check their host environment for supported values.

PMIX_COLLECTIVE_ALGO_REQD "pmix.calreqd" (bool)

If **true**, indicates that the requested choice of algorithm is mandatory.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Note that PMIx libraries may choose to implement an optimization for the case where only the calling process is involved in the fence operation by immediately returning **PMIX_OPERATION_SUCCEEDED** from the client’s call in lieu of passing the fence operation to a PMIx server. Fence operations involving more than just the calling process must be communicated to the PMIx server for proper execution of the included barrier behavior.

Similarly, fence operations that involve only processes that are clients of the same PMIx server may be resolved by that server without referral to its host environment as no inter-node coordination is required.

Description

Nonblocking **PMIx_Fence** routine. Note that the function will return an error if a **NULL** callback function is given.

Note that for scalability reasons, the default behavior for **PMIx_Fence_nb** is to not collect the data.

See the **PMIx_Fence** description for further details.

1 5.3 Publish and Lookup Data

2 The APIs defined in this section publish data from one client that can be later exchanged and looked
3 up by another client.

▼ **Advice to PMIx library implementers** ▼

4 PMIx libraries that support any of the functions in this section are required to support *all* of them.

▲

▼ **Advice to PMIx server hosts** ▼

5 Host environments that support any of the functions in this section are required to support *all* of
6 them.

▲

7 5.3.1 PMIx_Publish

8 **Summary**

9 Publish data for later access via [PMIx_Lookup](#) .

10 **Format**

PMIx v1.0

▼ **C** ▼

11 `pmix_status_t`

12 `PMIx_Publish(const pmix_info_t info[], size_t ninfo)`

▲ **C** ▲

13 **IN** `info`

14 Array of info structures (array of handles)

15 **IN** `ninfo`

16 Number of element in the *info* array (integer)

17 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

▼ **Required Attributes** ▼

18 PMIx libraries are not required to directly support any attributes for this function. However, any
19 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
20 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that
21 published the info.

▲

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

PMIX_PERSISTENCE "pmix.persist" (pmix_persistence_t)

Value for calls to **PMIx_Publish** .

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Publish the data in the *info* array for subsequent lookup. By default, the data will be published into the **PMIX_RANGE_SESSION** range and with **PMIX_PERSIST_APP** persistence. Changes to those values, and any additional directives, can be included in the **pmix_info_t** array. Attempts to access the data by processes outside of the provided data range will be rejected. The persistence parameter instructs the server as to how long the data is to be retained.

The blocking form will block until the server confirms that the data has been sent to the PMIx server and that it has obtained confirmation from its host SMS daemon that the data is ready to be looked up. Data is copied into the backing key-value data store, and therefore the *info* array can be released upon return from the blocking function call.

Advice to users

Publishing duplicate keys is permitted provided they are published to different ranges.

Advice to PMIx library implementers

Implementations should, to the best of their ability, detect duplicate keys being posted on the same data range and protect the user from unexpected behavior by returning the **PMIX_ERR_DUPLICATE_KEY** error.

1 5.3.2 PMIx_Publish_nb

2 Summary

3 Nonblocking [PMIx_Publish](#) routine.

4 Format

PMIx v1.0

```
5 pmix_status_t  
6 PMIx_Publish_nb(const pmix_info_t info[], size_t ninfo,  
7 pmix_op_cbfunc_t cbfunc, void *cbdata)
```

8 IN info

9 Array of info structures (array of handles)

10 IN ninfo

11 Number of element in the *info* array (integer)

12 IN cbfunc

13 Callback function [pmix_op_cbfunc_t](#) (function reference)

14 IN cbdata

15 Data to be passed to the callback function (memory reference)

16 Returns one of the following:

- 17 • [PMIX_SUCCESS](#) , indicating that the request is being processed by the host environment - result
18 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
19 function prior to returning from the API.
- 20 • [PMIX_OPERATION_SUCCEEDED](#) , indicating that the request was immediately processed and
21 returned *success* - the *cbfunc* will *not* be called
- 22 • a PMIx error constant indicating either an error in the input or that the request was immediately
23 processed and failed - the *cbfunc* will *not* be called

Required Attributes

24 PMIx libraries are not required to directly support any attributes for this function. However, any
25 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
26 *required* to add the [PMIX_USERID](#) and the [PMIX_GRPID](#) attributes of the client process that
27 published the info.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

PMIX_PERSISTENCE "pmix.persist" (pmix_persistence_t)

Value for calls to **PMIx_Publish**.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Nonblocking **PMIx_Publish** routine. The non-blocking form will return immediately, executing the callback when the PMIx server receives confirmation from its host SMS daemon.

Note that the function will return an error if a **NULL** callback function is given, and that the *info* array must be maintained until the callback is provided.

5.3.3 PMIx_Lookup

Summary

Lookup information published by this or another process with **PMIx_Publish** or **PMIx_Publish_nb**.

1
PMIx v1.0

Format

```

pmix_status_t
PMIx_Lookup(pmix_pdata_t data[], size_t ndata,
            const pmix_info_t info[], size_t ninfo)

```

INOUT data

Array of publishable data structures (array of handles)

IN ndata

Number of elements in the *data* array (integer)

IN info

Array of info structures (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is requesting the info.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid "hangs" due to programming errors that prevent the target process from ever exposing its data.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

PMIX_WAIT "pmix.wait" (int)

Caller requests that the PMIx server wait until at least the specified number of values are found (0 indicates all and is the default).

Advice to PMIx library implementers

We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not created.

Description

Lookup information published by this or another process. By default, the search will be conducted across the `PMIX_RANGE_SESSION` range. Changes to the range, and any additional directives, can be provided in the `pmix_info_t` array. Data is returned provided the following conditions are met:

- the requesting process resides within the range specified by the publisher. For example, data published to `PMIX_RANGE_LOCAL` can only be discovered by a process executing on the same node
- the provided key matches the published key within that data range
- the data was published by a process with corresponding user and/or group IDs as the one looking up the data. There currently is no option to override this behavior - such an option may become available later via an appropriate `pmix_info_t` directive.

The *data* parameter consists of an array of `pmix_pdata_t` struct with the keys specifying the requested information. Data will be returned for each key in the associated *value* struct. Any key that cannot be found will return with a data type of `PMIX_UNDEF`. The function will return `PMIX_SUCCESS` if any values can be found, so the caller must check each data element to ensure it was returned.

The *proc* field in each `pmix_pdata_t` struct will contain the namespace/rank of the process that published the data.

Advice to users

Although this is a blocking function, it will not wait by default for the requested data to be published. Instead, it will block for the time required by the server to lookup its current data and return any found items. Thus, the caller is responsible for ensuring that data is published prior to executing a lookup, using `PMIX_WAIT` to instruct the server to wait for the data to be published, or for retrying until the requested data is found.

1 5.3.4 PMIx_Lookup_nb

2 Summary

3 Nonblocking version of [PMIx_Lookup](#) .

4 Format

PMIx v1.0

C

5 `pmix_status_t`

```
6 PMIx_Lookup_nb(char **keys,  
7           const pmix_info_t info[], size_t ninfo,  
8           pmix_lookup_cbfunc_t cbfunc, void *cbdata)
```

C

9 **IN keys**

10 Array to be provided to the callback (array of strings)

11 **IN info**

12 Array of info structures (array of handles)

13 **IN ninfo**

14 Number of element in the *info* array (integer)

15 **IN cbfunc**

16 Callback function (handle)

17 **IN cbdata**

18 Callback data to be provided to the callback function (pointer)

19 Returns one of the following:

- 20 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
21 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
22 function prior to returning from the API.
- 23 • a PMIx error constant indicating an error in the input - the *cbfunc* will *not* be called

Required Attributes

24 PMIx libraries are not required to directly support any attributes for this function. However, any
25 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
26 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is
27 requesting the info.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

PMIX_WAIT "pmix.wait" (int)

Caller requests that the PMIx server wait until at least the specified number of values are found (0 indicates all and is the default).

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking form of the **PMIx_Lookup** function. Data for the provided NULL-terminated *keys* array will be returned in the provided callback function. As with **PMIx_Lookup**, the default behavior is to not wait for data to be published. The *info* array can be used to modify the behavior as previously described by **PMIx_Lookup**. Both the *info* and *keys* arrays must be maintained until the callback is provided.

5.3.5 PMIx_Unpublish

Summary

Unpublish data posted by this process using the given keys.

1
PMIx v1.0

Format

C

```
2 pmix_status_t
3 PMIx_Unpublish(char **keys,
4                 const pmix_info_t info[], size_t ninfo)
```

C

- 5 **IN info**
- 6 Array of info structures (array of handles)
- 7 **IN ninfo**
- 8 Number of element in the *info* array (integer)

9 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Required Attributes

10 PMIx libraries are not required to directly support any attributes for this function. However, any
11 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
12 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is
13 requesting the operation.



Optional Attributes

14 The following attributes are optional for host environments that support this operation:

15 **PMIX_TIMEOUT** "pmix.timeout" (**int**)
16 Time in seconds before the specified operation should time out (0 indicating infinite) in
17 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
18 the target process from ever exposing its data.

19 **PMIX_RANGE** "pmix.range" (**pmix_data_range_t**)
20 Value for calls to publish/lookup/unpublish or for monitoring event notifications.



Advice to PMIx library implementers

21 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
22 environment due to race condition considerations between completion of the operation versus
23 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
24 directly in the PMIx server library must take care to resolve the race condition and should avoid
25 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
26 created.



1 **Description**
 2 Unpublish data posted by this process using the given *keys*. The function will block until the data
 3 has been removed by the server (i.e., it is safe to publish that key again). A value of **NULL** for the
 4 *keys* parameter instructs the server to remove all data published by this process.
 5 By default, the range is assumed to be **PMIX_RANGE_SESSION** . Changes to the range, and any
 6 additional directives, can be provided in the *info* array.

7 5.3.6 PMIx_Unpublish_nb

8 **Summary**
 9 Nonblocking version of **PMIx_Unpublish** .

10 **Format**

PMIx v1.0 

```

11 pmix_status_t
12 PMIx_Unpublish_nb(char **keys,
13                  const pmix_info_t info[], size_t ninfo,
14                  pmix_op_cbfunc_t cbfunc, void *cbdata)
  
```

- 15 **IN keys**
 16 (array of strings)
- 17 **IN info**
 18 Array of info structures (array of handles)
- 19 **IN ninfo**
 20 Number of element in the *info* array (integer)
- 21 **IN cbfunc**
 22 Callback function **pmix_op_cbfunc_t** (function reference)
- 23 **IN cbdata**
 24 Data to be passed to the callback function (memory reference)

- 25 Returns one of the following:
- 26 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
 27 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
 28 function prior to returning from the API.
 - 29 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
 30 returned *success* - the *cbfunc* will *not* be called
 - 31 • a PMIx error constant indicating either an error in the input or that the request was immediately
 32 processed and failed - the *cbfunc* will *not* be called

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is requesting the operation.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking form of the **PMIx_Unpublish** function. The callback function will be executed once the server confirms removal of the specified data. The *info* array must be maintained until the callback is provided.

CHAPTER 6

Process Management

1 This chapter defines functionality used by clients to create and destroy/abort processes in the PMIx
2 universe.

3 6.1 Abort

4 PMIx provides a dedicated API by which an application can request that specified processes be
5 aborted by the system.

6 6.1.1 PMIx_Abort

7 Summary

8 Abort the specified processes

9 Format

PMIx v1.0

C

```
10 pmix_status_t  
11 PMIx_Abort(int status, const char msg[],  
12           pmix_proc_t procs[], size_t nprocs)
```

C

13 **IN** **status**

14 Error code to return to invoking environment (integer)

15 **IN** **msg**

16 String message to be returned to user (string)

17 **IN** **procs**

18 Array of `pmix_proc_t` structures (array of handles)

19 **IN** **nprocs**

20 Number of elements in the *procs* array (integer)

21 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Request that the host resource manager print the provided message and abort the provided array of *procs*. A Unix or POSIX environment should handle the provided status as a return error code from the main program that launched the application. A **NULL** for the *procs* array indicates that all processes in the caller's namespace are to be aborted, including itself. Passing a **NULL** *msg* parameter is allowed.

Advice to users

The response to this request is somewhat dependent on the specific resource manager and its configuration (e.g., some resource managers will not abort the application if the provided status is zero unless specifically configured to do so, and some cannot abort subsets of processes in an application), and thus lies outside the control of PMIx itself. However, the PMIx client library shall inform the RM of the request that the specified *procs* be aborted, regardless of the value of the provided status.

Note that race conditions caused by multiple processes calling **PMIx_Abort** are left to the server implementation to resolve with regard to which status is returned and what messages (if any) are printed.

6.2 Process Creation

The **PMIx_Spawn** commands spawn new processes and/or applications in the PMIx universe. This may include requests to extend the existing resource allocation or obtain a new one, depending upon provided and supported attributes.

6.2.1 PMIx_Spawn

Summary

Spawn a new job.

1
PMIx v1.0

Format

C

```
2 pmix_status_t
3 PMIx_Spawn(const pmix_info_t job_info[], size_t ninfo,
4             const pmix_app_t apps[], size_t napps,
5             char nspace[])
```

C

6 **IN** `job_info`
7 Array of info structures (array of handles)
8 **IN** `ninfo`
9 Number of elements in the `job_info` array (integer)
10 **IN** `apps`
11 Array of `pmix_app_t` structures (array of handles)
12 **IN** `napps`
13 Number of elements in the `apps` array (integer)
14 **OUT** `nspace`
15 Namespace of the new job (string)

16 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Required Attributes

17 PMIx libraries are not required to directly support any attributes for this function. However, any
18 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
19 required to add the following attributes to those provided before passing the request to the host:

- 20 **PMIX_SPAWNED** "pmix.spawned" (bool)
21 `true` if this process resulted from a call to `PMIx_Spawn`.
- 22 **PMIX_PARENT_ID** "pmix.parent" (`pmix_proc_t`)
23 Process identifier of the parent process of the calling process.
- 24 **PMIX_REQUESTOR_IS_CLIENT** "pmix.req.client" (bool)
25 The requesting process is a PMIx client.
- 26 **PMIX_REQUESTOR_IS_TOOL** "pmix.req.tool" (bool)
27 The requesting process is a PMIx tool.

28
29 Host environments that implement support for `PMIx_Spawn` are required to pass the
30 `PMIX_SPAWNED` and `PMIX_PARENT_ID` attributes to all PMIx servers launching new child
31 processes so those values can be returned to clients upon connection to the PMIx server. In
32 addition, they are required to support the following attributes when present in either the `job_info` or
33 the `info` array of an element of the `apps` array:

- 34 **PMIX_WDIR** "pmix.wdir" (`char*`)

1 Working directory for spawned processes.

2 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

3 Set the application's current working directory to the session working directory assigned by
4 the RM - when accessed using **PMIx_Get** , use the **PMIX_RANK_WILDCARD** value for
5 the rank to discover the session working directory assigned to the provided namespace

6 **PMIX_PREFIX** "pmix.prefix" (char*)


7 Prefix to use for starting spawned processes.

8 **PMIX_HOST** "pmix.host" (char*)

9 Comma-delimited list of hosts to use for spawned processes.

10 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

11 Hostfile to use for spawned processes.



Optional Attributes

12 The following attributes are optional for host environments that support this operation:

13 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

14 Hostfile listing hosts to add to existing allocation.

15 **PMIX_ADD_HOST** "pmix.addhost" (char*)

16 Comma-delimited list of hosts to add to the allocation.

17 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

18 Preload binaries onto nodes.

19 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

20 Comma-delimited list of files to pre-position on nodes.

21 **PMIX_PERSONALITY** "pmix.pers" (char*)

22 Name of personality to use.

23 **PMIX_MAPPER** "pmix.mapper" (char*)

24 Mapping mechanism to use for placing spawned processes - when accessed using
25 **PMIx_Get** , use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping
26 mechanism used for the provided namespace.

27 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)

28 Display process mapping upon spawn.

29 **PMIX_PPR** "pmix.ppr" (char*)

30 Number of processes to spawn on each identified resource.

31 **PMIX_MAPBY** "pmix.mapby" (char*)

1 Process mapping policy - when accessed using `PMIx_Get` , use the
2 `PMIX_RANK_WILDCARD` value for the rank to discover the mapping policy used for the
3 provided namespace

4 `PMIX_RANKBY` "pmix.rankby" (char*)
5 Process ranking policy - when accessed using `PMIx_Get` , use the
6 `PMIX_RANK_WILDCARD` value for the rank to discover the ranking algorithm used for the
7 provided namespace

8 `PMIX_BINDTO` "pmix.bindto" (char*)
9 Process binding policy - when accessed using `PMIx_Get` , use the
10 `PMIX_RANK_WILDCARD` value for the rank to discover the binding policy used for the
11 provided namespace

12 `PMIX_NON_PMI` "pmix.nonpmi" (bool)
13 Spawned processes will not call `PMIx_Init` .

14 `PMIX_STDIN_TGT` "pmix.stdin" (uint32_t)
15 Spawned process rank that is to receive `stdin`.

16 `PMIX_FWD_STDIN` "pmix.fwd.stdin" (bool)
17 Forward this process's `stdin` to the designated process.

18 `PMIX_FWD_STDOUT` "pmix.fwd.stdout" (bool)
19 Forward `stdout` from spawned processes to this process.

20 `PMIX_FWD_STDERR` "pmix.fwd.stderr" (bool)
21 Forward `stderr` from spawned processes to this process.

22 `PMIX_DEBUGGER_DAEMONS` "pmix.debugger" (bool)
23 Spawned application consists of debugger daemons.

24 `PMIX_TAG_OUTPUT` "pmix.tagout" (bool)
25 Tag application output with the identity of the source process.

26 `PMIX_TIMESTAMP_OUTPUT` "pmix.tsout" (bool)
27 Timestamp output from applications.

28 `PMIX_MERGE_STDERR_STDOUT` "pmix.mergeerrout" (bool)
29 Merge `stdout` and `stderr` streams from application processes.

30 `PMIX_OUTPUT_TO_FILE` "pmix.outfile" (char*)
31 Output application output to the specified file.

32 `PMIX_INDEX_ARGV` "pmix.indxargv" (bool)
33 Mark the `argv` with the rank of the process.

34 `PMIX_CPUS_PER_PROC` "pmix.cpusperproc" (uint32_t)

1 Number of cpus to assign to each rank - when accessed using `PMIx_Get` , use the
2 `PMIX_RANK_WILDCARD` value for the rank to discover the cpus/process assigned to the
3 provided namespace

4 `PMIX_NO_PROCS_ON_HEAD` "pmix.nolocal" (bool)

5 Do not place processes on the head node.

6 `PMIX_NO_OVERSUBSCRIBE` "pmix.noover" (bool)

7 Do not oversubscribe the cpus.

8 `PMIX_REPORT_BINDINGS` "pmix.repbinding" (bool)

9 Report bindings of the individual processes.

10 `PMIX_CPU_LIST` "pmix.cpulist" (char*)

11 List of cpus to use for this job - when accessed using `PMIx_Get` , use the
12 `PMIX_RANK_WILDCARD` value for the rank to discover the cpu list used for the provided
13 namespace

14 `PMIX_JOB_RECOVERABLE` "pmix.recover" (bool)

15 Application supports recoverable operations.

16 `PMIX_JOB_CONTINUOUS` "pmix.continuous" (bool)

17 Application is continuous, all failed processes should be immediately restarted.

18 `PMIX_MAX_RESTARTS` "pmix.maxrestarts" (uint32_t)

19 Maximum number of times to restart a job - when accessed using `PMIx_Get` , use the
20 `PMIX_RANK_WILDCARD` value for the rank to discover the max restarts for the provided
21 namespace

22 `PMIX_NOTIFY_COMPLETION` "pmix.notecomp" (bool)

23 Notify the parent process upon termination of child job.



24 **Description**

25 Spawn a new job. The assigned namespace of the spawned applications is returned in the *nspc*
26 parameter. A **NULL** value in that location indicates that the caller doesn't wish to have the
27 namespace returned. The *nspc* array must be at least of size one more than `PMIX_MAX_NSPLEN` .

28 By default, the spawned processes will be PMIx “connected” to the parent process upon successful
29 launch (see `PMIx_Connect` description for details). Note that this only means that (a) the parent
30 process will be given a copy of the new job's information so it can query job-level info without
31 incurring any communication penalties, (b) newly spawned child processes will receive a copy of
32 the parent processes job-level info, and (c) both the parent process and members of the child job
33 will receive notification of errors from processes in their combined assemblage.

Advice to users

Behavior of individual resource managers may differ, but it is expected that failure of any application process to start will result in termination/cleanup of all processes in the newly spawned job and return of an error code to the caller.

6.2.2 PMIx_Spawn_nb

Summary

Nonblocking version of the `PMIx_Spawn` routine.

Format

PMIx v1.0

C

```
pmix_status_t
PMIx_Spawn_nb(const pmix_info_t job_info[], size_t ninfo,
              const pmix_app_t apps[], size_t napps,
              pmix_spawn_cbfunc_t cbfunc, void *cbdata)
```

C

IN `job_info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the `job_info` array (integer)

IN `apps`

Array of `pmix_app_t` structures (array of handles)

IN `cbfunc`

Callback function `pmix_spawn_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- `PMIX_SUCCESS`, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API.
- a PMIx error constant indicating an error in the request - the `cbfunc` will *not* be called

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is required to add the following attributes to those provided before passing the request to the host:

PMIX_SPAWNED "pmix.spawned" (bool)

true if this process resulted from a call to [PMIx_Spawn](#).

PMIX_PARENT_ID "pmix.parent" (pmix_proc_t)

Process identifier of the parent process of the calling process.

PMIX_REQUESTOR_IS_CLIENT "pmix.req.client" (bool)

The requesting process is a PMIx client.

PMIX_REQUESTOR_IS_TOOL "pmix.req.tool" (bool)

The requesting process is a PMIx tool.

Host environments that implement support for [PMIx_Spawn](#) are required to pass the [PMIX_SPAWNED](#) and [PMIX_PARENT_ID](#) attributes to all PMIx servers launching new child processes so those values can be returned to clients upon connection to the PMIx server. In addition, they are required to support the following attributes when present in either the *job_info* or the *info* array of an element of the *apps* array:

PMIX_WDIR "pmix.wdir" (char*)

Working directory for spawned processes.

PMIX_SET_SESSION_CWD "pmix.ssn cwd" (bool)

Set the application's current working directory to the session working directory assigned by the RM - when accessed using [PMIx_Get](#), use the [PMIX_RANK_WILDCARD](#) value for the rank to discover the session working directory assigned to the provided namespace

PMIX_PREFIX "pmix.prefix" (char*)

Prefix to use for starting spawned processes.

PMIX_HOST "pmix.host" (char*)

Comma-delimited list of hosts to use for spawned processes.

PMIX_HOSTFILE "pmix.hostfile" (char*)

Hostfile to use for spawned processes.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_ADD_HOSTFILE "pmix.addhostfile" (char*)

Hostfile listing hosts to add to existing allocation.

PMIX_ADD_HOST "pmix.addhost" (char*)

Comma-delimited list of hosts to add to the allocation.

PMIX_PRELOAD_BIN "pmix.preloadbin" (bool)

Preload binaries onto nodes.

PMIX_PRELOAD_FILES "pmix.preloadfiles" (char*)

Comma-delimited list of files to pre-position on nodes.

PMIX_PERSONALITY "pmix.pers" (char*)

Name of personality to use.

PMIX_MAPPER "pmix.mapper" (char*)

Mapping mechanism to use for placing spawned processes - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping mechanism used for the provided namespace.

PMIX_DISPLAY_MAP "pmix.dispmap" (bool)

Display process mapping upon spawn.

PMIX_PPR "pmix.ppr" (char*)

Number of processes to spawn on each identified resource.

PMIX_MAPBY "pmix.mapby" (char*)

Process mapping policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the provided namespace

PMIX_RANKBY "pmix.rankby" (char*)

Process ranking policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the provided namespace

PMIX_BINDTO "pmix.bindto" (char*)

Process binding policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the provided namespace

PMIX_NON_PMI "pmix.nonpmi" (bool)

Spawned processes will not call **PMIx_Init**.

PMIX_STDIN_TGT "pmix.stdin" (uint32_t)

Spawned process rank that is to receive **stdin**.

1 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (bool)
 2 Forward this process's **stdin** to the designated process.

3 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)
 4 Forward **stdout** from spawned processes to this process.

5 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)
 6 Forward **stderr** from spawned processes to this process.

7 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)
 8 Spawned application consists of debugger daemons.

9 **PMIX_TAG_OUTPUT** "pmix.tagout" (bool)
 10 Tag application output with the identity of the source process.

11 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
 12 Timestamp output from applications.

13 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)
 14 Merge **stdout** and **stderr** streams from application processes.

15 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)
 16 Output application output to the specified file.

17 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)
 18 Mark the **argv** with the rank of the process.

19 **PMIX_CPUS_PER_PROC** "pmix.cpusperproc" (uint32_t)
 20 Number of cpus to assign to each rank - when accessed using **PMIx_Get** , use the
 21 **PMIX_RANK_WILDCARD** value for the rank to discover the cpus/process assigned to the
 22 provided namespace

23 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)
 24 Do not place processes on the head node.

25 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)
 26 Do not oversubscribe the cpus.

27 **PMIX_REPORT_BINDINGS** "pmix.repbinding" (bool)
 28 Report bindings of the individual processes.

29 **PMIX_CPU_LIST** "pmix.cpulists" (char*)
 30 List of cpus to use for this job - when accessed using **PMIx_Get** , use the
 31 **PMIX_RANK_WILDCARD** value for the rank to discover the cpu list used for the provided
 32 namespace

33 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)
 34 Application supports recoverable operations.

35 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)
 36 Application is continuous, all failed processes should be immediately restarted.

1 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)
2 Maximum number of times to restart a job - when accessed using **PMIx_Get** , use the
3 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
4 namespace



5 **Description**

6 Nonblocking version of the **PMIx_Spawn** routine. The provided callback function will be
7 executed upon successful start of *all* specified application processes.

▼ Advice to users ▼

8 Behavior of individual resource managers may differ, but it is expected that failure of any
9 application process to start will result in termination/cleanup of all processes in the newly spawned
10 job and return of an error code to the caller.



11 **6.3 Connecting and Disconnecting Processes**

12 This section defines functions to connect and disconnect processes in two or more separate PMIx
13 namespaces. The PMIx definition of *connected* solely implies that the host environment should
14 treat the failure of any process in the assemblage as a reportable event, taking action on the
15 assemblage as if it were a single application. For example, if the environment defaults (in the
16 absence of any application directives) to terminating an application upon failure of any process in
17 that application, then the environment should terminate all processes in the connected assemblage
18 upon failure of any member.

▼ Advice to PMIx server hosts ▼

19 The host environment may choose to assign a new namespace to the connected assemblage and/or
20 assign new ranks for its members for its own internal tracking purposes. However, it is not required
21 to communicate such assignments to the participants (e.g., in response to an appropriate call to
22 **PMIx_Query_info_nb**). The host environment is required to generate a
23 **PMIX_ERR_INVALID_TERMINATION** event should any process in the assemblage terminate or
24 call **PMIx_Finalize** without first *disconnecting* from the assemblage.

25 The *connect* operation does not require the exchange of job-level information nor the inclusion of
26 information posted by participating processes via **PMIx_Put** . Indeed, the callback function
27 utilized in **pmix_server_connect_fn_t** cannot pass information back into the PMIx server
28 library. However, host environments are advised that collecting such information at the
29 participating daemons represents an optimization opportunity as participating processes are likely
30 to request such information after the connect operation completes.



Advice to users

1 Attempting to *connect* processes solely within the same namespace is essentially a *no-op* operation.
2 While not explicitly prohibited, users are advised that a PMIx implementation or host environment
3 may return an error in such cases.

4 Neither the PMIx implementation nor host environment are required to provide any tracking
5 support for the assemblage. Thus, the application is responsible for maintaining the membership
6 list of the assemblage.

6.3.1 PMIx_Connect

Summary

Connect namespaces.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_Connect(const pmix_proc_t procs[], size_t nprocs,  
             const pmix_info_t info[], size_t ninfo)
```

- IN procs**
Array of proc structures (array of handles)
- IN nprocs**
Number of elements in the *procs* array (integer)
- IN info**
Array of info structures (array of handles)
- IN ninfo**
Number of elements in the *info* array (integer)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_COLLECTIVE_ALGO "pmix.calgo" (char*)

Comma-delimited list of algorithms to use for the collective operation. PMIx does not impose any requirements on a host environment’s collective algorithms. Thus, the acceptable values for this attribute will be environment-dependent - users are encouraged to check their host environment for supported values.

PMIX_COLLECTIVE_ALGO_REQD "pmix.calreqd" (bool)

If **true**, indicates that the requested choice of algorithm is mandatory.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The function will return once all processes identified in *procs* have called either **PMIx_Connect** or its non-blocking version, *and* the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes.

Advice to users

All processes engaged in a given **PMIx_Connect** operation must provide the identical *procs* array as ordering of entries in the array and the method by which those processes are identified (e.g., use of **PMIX_RANK_WILDCARD** versus listing the individual processes) *may* impact the host environment's algorithm for uniquely identifying an operation.

Advice to PMIx library implementers

PMIx_Connect and its non-blocking form are both *collective* operations. Accordingly, the PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

Processes that combine via **PMIx_Connect** must call **PMIx_Disconnect** prior to finalizing and/or terminating - any process in the assemblage failing to meet this requirement will cause a **PMIX_ERR_INVALID_TERMINATION** event to be generated.

A process can only engage in one connect operation involving the identical *procs* array at a time. However, a process can be simultaneously engaged in multiple connect operations, each involving a different *procs* array.

As in the case of the **PMIx_Fence** operation, the *info* array can be used to pass user-level directives regarding the algorithm to be used for any collective operation involved in the operation, timeout constraints, and other options available from the host RM.

6.3.2 **PMIx_Connect_nb**

Summary

Nonblocking **PMIx_Connect_nb** routine.

1
PMIx v1.0

Format

C

```

2  pmix_status_t
3  PMIx_Connect_nb(const pmix_proc_t procs[], size_t nprocs,
4                  const pmix_info_t info[], size_t ninfo,
5                  pmix_op_cbfunc_t cbfunc, void *cbdata)

```

C

- 6 **IN procs**
7 Array of proc structures (array of handles)
- 8 **IN nprocs**
9 Number of elements in the *procs* array (integer)
- 10 **IN info**
11 Array of info structures (array of handles)
- 12 **IN ninfo**
13 Number of element in the *info* array (integer)
- 14 **IN cbfunc**
15 Callback function `pmix_op_cbfunc_t` (function reference)
- 16 **IN cbdata**
17 Data to be passed to the callback function (memory reference)

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
21 function prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will *not* be called
- 24 • a PMIx error constant indicating either an error in the input or that the request was immediately
25 processed and failed - the *cbfunc* will *not* be called

Required Attributes

26 PMIx libraries are not required to directly support any attributes for this function. However, any
27 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_COLLECTIVE_ALGO "pmix.calgo" (char*)

Comma-delimited list of algorithms to use for the collective operation. PMIx does not impose any requirements on a host environment’s collective algorithms. Thus, the acceptable values for this attribute will be environment-dependent - users are encouraged to check their host environment for supported values.

PMIX_COLLECTIVE_ALGO_REQD "pmix.calreqd" (bool)

If **true**, indicates that the requested choice of algorithm is mandatory.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Nonblocking version of **PMIx_Connect**. The callback function is called once all processes identified in *procs* have called either **PMIx_Connect** or its non-blocking version, *and* the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes. See the advice provided in the description for **PMIx_Connect** for more information.

6.3.3 PMIx_Disconnect

Summary

Disconnect a previously connected set of processes.

1
PMIx v1.0

Format

C

```
2 pmix_status_t  
3 PMIx_Disconnect(const pmix_proc_t procs[], size_t nprocs,  
4 const pmix_info_t info[], size_t ninfo);
```

C

- 5 **IN procs**
6 Array of proc structures (array of handles)
- 7 **IN nprocs**
8 Number of elements in the *procs* array (integer)
- 9 **IN info**
10 Array of info structures (array of handles)
- 11 **IN ninfo**
12 Number of element in the *info* array (integer)

13 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Required Attributes

14 PMIx libraries are not required to directly support any attributes for this function. However, any
15 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

16 The following attributes are optional for host environments that support this operation:

17 **PMIX_TIMEOUT** "pmix.timeout" (**int**)
 18 Time in seconds before the specified operation should time out (0 indicating infinite) in
 19 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
 20 the target process from ever exposing its data.

Advice to PMIx library implementers

21 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
 22 environment due to race condition considerations between completion of the operation versus
 23 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
 24 directly in the PMIx server library must take care to resolve the race condition and should avoid
 25 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
 26 created.

Description

Disconnect a previously connected set of processes. A `PMIX_ERR_INVALID_OPERATION` error will be returned if the specified set of *procs* was not previously *connected* via a call to `PMIx_Connect` or its non-blocking form. The function will return once all processes identified in *procs* have called either `PMIx_Disconnect` or its non-blocking version, *and* the host environment has completed any required supporting operations.

Advice to users

All processes engaged in a given `PMIx_Disconnect` operation must provide the identical *procs* array as ordering of entries in the array and the method by which those processes are identified (e.g., use of `PMIX_RANK_WILDCARD` versus listing the individual processes) *may* impact the host environment's algorithm for uniquely identifying an operation.

Advice to PMIx library implementers

`PMIx_Disconnect` and its non-blocking form are both *collective* operations. Accordingly, the PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

A process can only engage in one disconnect operation involving the identical *procs* array at a time. However, a process can be simultaneously engaged in multiple disconnect operations, each involving a different *procs* array.

As in the case of the `PMIx_Fence` operation, the *info* array can be used to pass user-level directives regarding the algorithm to be used for any collective operation involved in the operation, timeout constraints, and other options available from the host RM.

6.3.4 `PMIx_Disconnect_nb`

Summary

Nonblocking `PMIx_Disconnect` routine.

1
PMIx v1.0

Format

C

```

2 pmix_status_t
3 PMIx_Disconnect_nb(const pmix_proc_t procs[], size_t nprocs,
4                   const pmix_info_t info[], size_t ninfo,
5                   pmix_op_cbfunc_t cbfunc, void *cbdata);

```

C

- 6 **IN procs**
7 Array of proc structures (array of handles)
- 8 **IN nprocs**
9 Number of elements in the *procs* array (integer)
- 10 **IN info**
11 Array of info structures (array of handles)
- 12 **IN ninfo**
13 Number of element in the *info* array (integer)
- 14 **IN cbfunc**
15 Callback function `pmix_op_cbfunc_t` (function reference)
- 16 **IN cbdata**
17 Data to be passed to the callback function (memory reference)

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
21 function prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will *not* be called
- 24 • a PMIx error constant indicating either an error in the input or that the request was immediately
25 processed and failed - the *cbfunc* will *not* be called

Required Attributes

26 PMIx libraries are not required to directly support any attributes for this function. However, any
27 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

28 The following attributes are optional for host environments that support this operation:

- 29 **PMIX_TIMEOUT** "pmix.timeout" (int)
30 Time in seconds before the specified operation should time out (0 indicating infinite) in
31 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
32 the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not created.

Description

Nonblocking `PMIx_Disconnect` routine. The callback function is called once all processes identified in *procs* have called either `PMIx_Disconnect_nb` or its blocking version, *and* the host environment has completed any required supporting operations. See the advice provided in the description for `PMIx_Disconnect` for more information.

6.4 IO Forwarding

This section defines functions by which tools (e.g., debuggers) can request forwarding of input/output to/from other processes. The term “tool” widely refers to non-computational programs executed by the user or system administrator to monitor or control a principal computational program. Tools almost always interact with either the host environment, user applications, or both to perform administrative and support functions. For example, a debugger tool might be used to remotely control the processes of a parallel application, monitoring their behavior on a step-by-step basis.

Underlying the operation of many tools is a common need to forward stdin from the tool to targeted processes, and to return stdout/stderr from those processes for display on the user’s console. Historically, each tool developer was responsible for creating their own IO forwarding subsystem. However, with the introduction of PMIx as a standard mechanism for interacting between applications and the host environment, it has become possible to relieve tool developers of this burden.

Advice to PMIx server hosts

The responsibility of the host environment in forwarding of IO falls into the following areas:

- Capturing output from specified child processes
- Forwarding that output to the host of the PMIx server library that requested it
- Delivering that payload to the PMIx server library via the `PMIx_server_IOF_deliver` API for final dispatch

It is the responsibility of the PMIx library to buffer, format, and deliver the payload to the requesting client.

Advice to users

The forwarding of IO via PMIx requires that both the host environment and the tool support PMIx, but does not impose any similar requirements on the application itself.

6.4.1 PMIx_IOF_pull

Summary

Register to receive output forwarded from a set of remote processes.

Format

PMIx v3.0

C

`pmix_status_t`

```
PMIx_IOF_pull(const pmix_proc_t procs[], size_t nprocs,  
              const pmix_info_t directives[], size_t ndirs,  
              pmix_iof_channel_t channel, pmix_iof_cbfunc_t cbfunc,  
              pmix_hdlr_reg_cbfunc_t regcbfunc, void *regcbdata)
```

C

IN `procs`

Array of proc structures identifying desired source processes (array of handles)

IN `nprocs`

Number of elements in the *procs* array (integer)

IN `directives`

Array of `pmix_info_t` structures (array of handles)

IN `ndirs`

Number of elements in the *directives* array (integer)

IN `channel`

Bitmask of IO channels included in the request (`pmix_iof_channel_t`)

IN `cbfunc`

Callback function for delivering relevant output (`pmix_iof_cbfunc_t` function reference)

IN `regcbfunc`

Function to be called when registration is completed (`pmix_hdlr_reg_cbfunc_t` function reference)

IN `regcbdata`

Data to be passed to the *regcbfunc* callback function (memory reference)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant. In the event the function returns an error, the *regcbfunc* will *not* be called.

Required Attributes

The following attributes are required for PMIx libraries that support IO forwarding:

PMIX_IOF_CACHE_SIZE "pmix.iof.csize" (uint32_t)

The requested size of the server cache in bytes for each specified channel. By default, the server is allowed (but not required) to drop all bytes received beyond the max size.

PMIX_IOF_DROP_OLDEST "pmix.iof.old" (bool)

In an overflow situation, drop the oldest bytes to make room in the cache.

PMIX_IOF_DROP_NEWEST "pmix.iof.new" (bool)

In an overflow situation, drop any new bytes received until room becomes available in the cache (default).

Optional Attributes

The following attributes are optional for PMIx libraries that support IO forwarding:

PMIX_IOF_BUFFERING_SIZE "pmix.iof.bsize" (uint32_t)

Controls grouping of IO on the specified channel(s) to avoid being called every time a bit of IO arrives. The library will execute the callback whenever the specified number of bytes becomes available. Any remaining buffered data will be “flushed” upon call to deregister the respective channel.

PMIX_IOF_BUFFERING_TIME "pmix.iof.btime" (uint32_t)

Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering size, this prevents IO from being held indefinitely while waiting for another payload to arrive.

PMIX_IOF_TAG_OUTPUT "pmix.iof.tag" (bool)

Tag output with the channel it comes from.

PMIX_IOF_TIMESTAMP_OUTPUT "pmix.iof.ts" (bool)

Timestamp output

PMIX_IOF_XML_OUTPUT "pmix.iof.xml" (bool)

Format output in XML

Description

Register to receive output forwarded from a set of remote processes.

Advice to users

Providing a **NULL** function pointer for the *cbfunc* parameter will cause output for the indicated channels to be written to their corresponding stdout/stderr file descriptors. Use of **PMIX_RANK_WILDCARD** to specify all processes in a given namespace is supported but should be used carefully due to bandwidth considerations.

6.4.2 PMIx_IOF_deregister

Summary

Deregister from output forwarded from a set of remote processes.

Format

PMIx v3.0

```
pmix_status_t
PMIx_IOF_deregister(size_t iofhdlr,
                    const pmix_info_t directives[], size_t ndirs,
                    pmix_op_cbfunc_t cbfunc, void *cbdata)
```

IN iofhdlr

Registration number returned from the **pmix_hdlr_reg_cbfunc_t** callback from the call to **PMIx_IOF_pull** (**size_t**)

IN directives

Array of **pmix_info_t** structures (array of handles)

IN ndirs

Number of elements in the *directives* array (integer)

IN cbfunc

Callback function to be called when deregistration has been completed. (function reference)

IN cbdata

Data to be passed to the *cbfunc* callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will *not* be called

Description

Deregister from output forwarded from a set of remote processes.

Advice to PMIx library implementers

Any currently buffered IO should be flushed upon receipt of a deregistration request. All received IO after receipt of the request shall be discarded.

6.4.3 PMIx_IOF_push

Summary

Push data collected locally (typically from stdin or a file) to stdin of the target recipients.

Format

PMIx v3.0

C

```
pmix_status_t
PMIx_IOF_push(const pmix_proc_t targets[], size_t ntargets,
              pmix_byte_object_t *bo,
              const pmix_info_t directives[], size_t ndirs,
              pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

IN targets

Array of proc structures identifying desired target processes (array of handles)

IN ntargets

Number of elements in the *targets* array (integer)

IN bo

Pointer to [pmix_byte_object_t](#) containing the payload to be delivered (handle)

IN directives

Array of [pmix_info_t](#) structures (array of handles)

IN ndirs

Number of elements in the *directives* array (integer)

IN directives

Array of [pmix_info_t](#) structures (array of handles)

IN cbfunc

Callback function to be called when operation has been completed. ([pmix_op_cbfunc_t](#) function reference)

IN cbdata

Data to be passed to the *cbfunc* callback function (memory reference)

Returns one of the following:

- 1 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
2 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
3 function prior to returning from the API.
- 4 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
5 returned *success* - the *cbfunc* will *not* be called
- 6 • a PMIx error constant indicating either an error in the input or that the request was immediately
7 processed and failed - the *cbfunc* will *not* be called

Required Attributes

8 The following attributes are required for PMIx libraries that support IO forwarding:

9 **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (**uint32_t**)

10 The requested size of the server cache in bytes for each specified channel. By default, the
11 server is allowed (but not required) to drop all bytes received beyond the max size.

12 **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (**bool**)

13 In an overflow situation, drop the oldest bytes to make room in the cache.

14 **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (**bool**)

15 In an overflow situation, drop any new bytes received until room becomes available in the
16 cache (default).

Optional Attributes

17 The following attributes are optional for PMIx libraries that support IO forwarding:

18 **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (**uint32_t**)

19 Controls grouping of IO on the specified channel(s) to avoid being called every time a bit of
20 IO arrives. The library will execute the callback whenever the specified number of bytes
21 becomes available. Any remaining buffered data will be “flushed” upon call to deregister the
22 respective channel.

23 **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (**uint32_t**)

24 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
25 size, this prevents IO from being held indefinitely while waiting for another payload to
26 arrive.

1
2
3
4
5
6

Description

Push data collected locally (typically from stdin or a file) to stdin of the target recipients.

Advice to users

Execution of the *cbfunc* callback function serves as notice that the PMIx library no longer requires the caller to maintain the *bo* data object - it does *not* indicate delivery of the payload to the targets. Use of **PMIX_RANK_WILDCARD** to specify all processes in a given namespace is supported but should be used carefully due to bandwidth considerations.

CHAPTER 7

Job Management and Reporting

1 The job management APIs provide an application with the ability to orchestrate its operation in
2 partnership with the SMS. Members of this category include the
3 [PMIx_Allocation_request_nb](#), [PMIx_Job_control_nb](#), and
4 [PMIx_Process_monitor_nb](#) APIs.

5 7.1 Query

6 As the level of interaction between applications and the host SMS grows, so too does the need for
7 the application to query the SMS regarding its capabilities and state information. PMIx provides a
8 generalized query interface for this purpose, along with a set of standardized attribute keys to
9 support a range of requests. This includes requests to determine the status of scheduling queues and
10 active allocations, the scope of API and attribute support offered by the SMS, namespaces of active
11 jobs, location and information about a job's processes, and information regarding available
12 resources.

13 An example use-case for the [PMIx_Query_info_nb](#) API is to ensure clean job completion.
14 Time-shared systems frequently impose maximum run times when assigning jobs to resource
15 allocations. To shut down gracefully, e.g., to write a checkpoint before termination, it is necessary
16 for an application to periodically query the resource manager for the time remaining in its
17 allocation. This is especially true on systems for which allocation times may be shortened or
18 lengthened from the original time limit. Many resource managers provide APIs to dynamically
19 obtain this information, but each API is specific to the resource manager.

20 PMIx supports this use-case by defining an attribute key ([PMIX_TIME_REMAINING](#)) that can be
21 used with the [PMIx_Query_info_nb](#) interface to obtain the number of seconds remaining in
22 the current job allocation. Note that one could alternatively use the
23 [PMIx_Register_event_handler](#) API to register for an event indicating incipient job
24 termination, and then use the [PMIx_Job_control_nb](#) API to request that the host SMS
25 generate an event a specified amount of time prior to reaching the maximum run time. PMIx
26 provides such alternate methods as a means of maximizing the probability of a host system
27 supporting at least one method by which the application can obtain the desired service.

28 The following APIs support query of various session and environment values.


29 7.1.1 [PMIx_Resolve_peers](#)

30 Summary

31 Obtain the array of processes within the specified namespace that are executing on a given node.

1 **Format**
PMIx v1.0 

```
2 pmix_status_t  
3 PMIx_Resolve_peers(const char *nodename,  
4                   const pmix_namespace_t nspace,  
5                   pmix_proc_t **procs, size_t *nprocs)
```



- 6 **IN** **nodename**
7 Name of the node to query (string)
- 8 **IN** **nspace**
9 namespace (string)
- 10 **OUT** **procs**
11 Array of process structures (array of handles)
- 12 **OUT** **nprocs**
13 Number of elements in the *procs* array (integer)

14 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.


15 **Description**
16 Given a *nodename*, return the array of processes within the specified *nspace* that are executing on
17 that node. If the *nspace* is **NULL**, then all processes on the node will be returned. If the specified
18 node does not currently host any processes, then the returned array will be **NULL**, and *nprocs* will
19 be 0. The caller is responsible for releasing the *procs* array when done with it. The
20 **PMIX_PROC_FREE** macro is provided for this purpose.

21 7.1.2 PMIx_Resolve_nodes

22 **Summary**
23 Return a list of nodes hosting processes within the given namespace.

24 **Format**
PMIx v1.0 

```
25 pmix_status_t  
26 PMIx_Resolve_nodes(const char *nnamespace, char **nodelist)
```



- 27 **IN** **nnamespace**
28 Namespace (string)
- 29 **OUT** **nodelist**
30 Comma-delimited list of nodenames (string)

31 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

1 **Description**
2 Given a *namespace*, return the list of nodes hosting processes within that namespace. The returned
3 string will contain a comma-delimited list of nodenames. The caller is responsible for releasing the
4 string when done with it.

5 7.1.3 PMIx_Query_info

6 **Summary**
7 Query information about the system in general.

8 **Format**

PMIx v4.0

C

```
9 pmix_status_t  
10 PMIx_Query_info(pmix_query_t queries[], size_t nqueries,  
11 pmix_info_t *info[], size_t *ninfo)
```

C

12 **IN queries**
13 Array of query structures (array of handles)
14 **IN nqueries**
15 Number of elements in the *queries* array (integer)
16 **INOUT info**
17 Address where a pointer to an array of **pmix_info_t** containing the results of the query
18 can be returned (memory reference)
19 **INOUT ninfo**
20 Address where the number of elements in *info* can be returned (handle)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS** All data has been returned
- 23 • **PMIX_ERR_NOT_FOUND** None of the requested data was available
- 24 • **PMIX_ERR_PARTIAL_SUCCESS** Some of the data has been returned
- 25 • **PMIX_ERR_NOT_SUPPORTED** The host RM does not support this function
- 26 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

27 PMIx libraries that support this API are required to support the following attributes:

28 **PMIX_QUERY_REFRESH_CACHE** "pmix.qry.rfsh" (bool)
29 Retrieve updated information from server.
30 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)

1 Return information about the specified session. If information about a session other than the
2 one containing the requesting process is desired, then the attribute array must contain a
3 **PMIX_SESSION_ID** attribute identifying the desired target.

4 **PMIX_JOB_INFO** "pmix.job.info" (bool)

5 Return information about the specified job or namespace. If information about a job or
6 namespace other than the one containing the requesting process is desired, then the attribute
7 array must contain a **PMIX_JOBID** or **PMIX_NAMESPACE** attribute identifying the desired
8 target. Similarly, if information is requested about a job or namespace in a session other than
9 the one containing the requesting process, then an attribute identifying the target session
10 must be provided.

11 **PMIX_APP_INFO** "pmix.app.info" (bool)

12 Return information about the specified application. If information about an application other
13 than the one containing the requesting process is desired, then the attribute array must
14 contain a **PMIX_APPNUM** attribute identifying the desired target. Similarly, if information is
15 requested about an application in a job or session other than the one containing the requesting
16 process, then attributes identifying the target job and/or session must be provided.

17 **PMIX_NODE_INFO** "pmix.node.info" (bool)

18 Return information about the specified node. If information about a node other than the one
19 containing the requesting process is desired, then the attribute array must contain either the
20 **PMIX_NODEID** or **PMIX_HOSTNAME** attribute identifying the desired target.

21 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)

22 Process identifier Specifies the process ID whose information is being requested - e.g., a
23 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
24 request is for information on a specific process.

25 **PMIX_NAMESPACE** "pmix.namespace" (char*)

26 Namespace of the job. Specifies the namespace of the process whose information is being
27 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
28 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
29 information on a specific process.

30 **PMIX_RANK** "pmix.rank" (pmix_rank_t)

31 Process rank within the job. Specifies the rank of the process whose information is being
32 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
33 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
34 information on a specific process.

35 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)

36 Query list of supported attributes for specified APIs

37 **PMIX_CLIENT_ATTRIBUTES** "pmix.client.attrs" (bool)

38 Request attributes supported by the PMIx client library

39 **PMIX_SERVER_ATTRIBUTES** "pmix.srvr.attrs" (bool)

1 Request attributes supported by the PMIx server library

2 **PMIX_HOST_ATTRIBUTES** "pmix.host.attrs" (bool)

3 Request attributes supported by the host environment

4 **PMIX_TOOL_ATTRIBUTES** "pmix.setup.env" (bool)

5 Request attributes supported by the PMIx tool library functions

6 Note that inclusion of the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
7 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
8 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on
9 multiple specific processes therefore requires submitting multiple **pmix_query_t** structures,
10 each referencing one process.

11 PMIx libraries are not required to directly support any other attributes for this function. However,
12 any provided attributes must be passed to the host SMS daemon for processing, and the PMIx
13 library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client
14 process making the request.

16 Host environments that support this operation are required to support the following attributes as
17 qualifiers to the request:

18 **PMIX_PROCID** "pmix.procid" (**pmix_proc_t**)

19 Process identifier Specifies the process ID whose information is being requested - e.g., a
20 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
21 request is for information on a specific process.

22 **PMIX_NAMESPACE** "pmix.namespace" (**char***)

23 Namespace of the job. Specifies the namespace of the process whose information is being
24 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
25 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
26 information on a specific process.

27 **PMIX_RANK** "pmix.rank" (**pmix_rank_t**)

28 Process rank within the job. Specifies the rank of the process whose information is being
29 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
30 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
31 information on a specific process.

32 Note that inclusion of the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
33 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
34 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on
35 multiple specific processes therefore requires submitting multiple **pmix_query_t** structures,
36 each referencing one process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_QUERY_NAMESPACES "pmix.qry.ns" (char*)

Request a comma-delimited list of active namespaces.

PMIX_QUERY_JOB_STATUS "pmix.qry.jst" (pmix_status_t)

Status of a specified, currently executing job.

PMIX_QUERY_QUEUE_LIST "pmix.qry.qlst" (char*)

Request a comma-delimited list of scheduler queues.

PMIX_QUERY_QUEUE_STATUS "pmix.qry.qst" (TBD)

Status of a specified scheduler queue.

PMIX_QUERY_PROC_TABLE "pmix.qry.phtable" (char*)

Input namespace of the job whose information is being requested returns (**pmix_data_array_t**) an array of **pmix_proc_info_t** .

PMIX_QUERY_LOCAL_PROC_TABLE "pmix.qry.lptable" (char*)

Input namespace of the job whose information is being requested returns (**pmix_data_array_t**) an array of **pmix_proc_info_t** for processes in job on same node.

PMIX_QUERY_SPAWN_SUPPORT "pmix.qry.spawn" (bool)

Return a comma-delimited list of supported spawn attributes.

PMIX_QUERY_DEBUG_SUPPORT "pmix.qry.debug" (bool)

Return a comma-delimited list of supported debug attributes.

PMIX_QUERY_MEMORY_USAGE "pmix.qry.mem" (bool)

Return information on memory usage for the processes indicated in the qualifiers.

PMIX_QUERY_REPORT_AVG "pmix.qry.avg" (bool)

Report only average values for sampled information.

PMIX_QUERY_REPORT_MINMAX "pmix.qry.minmax" (bool)

Report minimum and maximum values.

PMIX_QUERY_ALLOC_STATUS "pmix.query.alloc" (char*)

String identifier of the allocation whose status is being requested.

PMIX_TIME_REMAINING "pmix.time.remaining" (char*)

Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.

PMIX_SERVER_URI "pmix.srvr.uri" (char*)

URI of the PMIx server to be contacted. Requests the URI of the specified PMIx server's PMIx connection. Defaults to requesting the information for the local PMIx server.

1 **PMIX_PROC_URI** "pmix.puri" (char*)
2 URI containing contact information for a given process. Requests the URI of the specified
3 PMIx server's out-of-band connection. Defaults to requesting the information for the local
4 PMIx server.



5 **Description**

6 Query information about the system in general. This can include a list of active namespaces,
7 network topology, etc. Also can be used to query node-specific info such as the list of peers
8 executing on a given node. We assume that the host RM will exercise appropriate access control on
9 the information.

10 The returned *status* indicates if requested data was found or not. The returned array of
11 **pmix_info_t** will contain each key that was provided and the corresponding value that was
12 found. Requests for keys that are not found will return the key paired with a value of type
13 **PMIX_UNDEF**. The caller is responsible for releasing the returned array.

▼ **Advice to PMIx library implementers** ▼

14 Information returned from **PMIx_Query_info** shall be locally cached so that retrieval by
15 subsequent calls to **PMIx_Get**, **PMIx_Query_info**, or **PMIx_Query_info_nb** can
16 succeed with minimal overhead. The local cache shall be checked prior to querying the PMIx
17 server and/or the host environment. Queries that include the **PMIX_QUERY_REFRESH_CACHE**
18 attribute shall bypass the local cache and retrieve a new value for the query, refreshing the values in
19 the cache upon return.



20 **7.1.4 PMIx_Query_info_nb**

21 **Summary**

22 Query information about the system in general.

1
PMIx v2.0

Format

C

```
2 pmix_status_t
3 PMIx_Query_info_nb(pmix_query_t queries[], size_t nqueries,
4 pmix_info_cbfunc_t cbfunc, void *cbdata)
```

C

- 5 **IN queries**
- 6 Array of query structures (array of handles)
- 7 **IN nqueries**
- 8 Number of elements in the *queries* array (integer)
- 9 **IN cbfunc**
- 10 Callback function `pmix_info_cbfunc_t` (function reference)
- 11 **IN cbdata**
- 12 Data to be passed to the callback function (memory reference)

13 Returns one of the following:

- 14 • **PMIX_SUCCESS** indicating that the request has been accepted for processing and the provided
- 15 callback function will be executed upon completion of the operation. Note that the library must
- 16 not invoke the callback function prior to returning from the API.
- 17 • a non-zero PMIx error constant indicating a reason for the request to have been rejected. In this
- 18 case, the provided callback function will not be executed

19 If executed, the status returned in the provided callback function will be one of the following

20 constants:

- 21 • **PMIX_SUCCESS** All data has been returned
- 22 • **PMIX_ERR_NOT_FOUND** None of the requested data was available
- 23 • **PMIX_ERR_PARTIAL_SUCCESS** Some of the data has been returned
- 24 • **PMIX_ERR_NOT_SUPPORTED** The host RM does not support this function
- 25 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

26 PMIx libraries that support this API are required to support the following attributes:

- 27 **PMIX_QUERY_REFRESH_CACHE** "`pmix.qry.rfsh`" (bool)
- 28 Retrieve updated information from server.
- 29 **PMIX_SESSION_INFO** "`pmix.ssn.info`" (bool)
- 30 Return information about the specified session. If information about a session other than the
- 31 one containing the requesting process is desired, then the attribute array must contain a
- 32 **PMIX_SESSION_ID** attribute identifying the desired target.

1 **PMIX_JOB_INFO** "pmix.job.info" (bool)
2 Return information about the specified job or namespace. If information about a job or
3 namespace other than the one containing the requesting process is desired, then the attribute
4 array must contain a **PMIX_JOBID** or **PMIX_NAMESPACE** attribute identifying the desired
5 target. Similarly, if information is requested about a job or namespace in a session other than
6 the one containing the requesting process, then an attribute identifying the target session
7 must be provided.

8 **PMIX_APP_INFO** "pmix.app.info" (bool)
9 Return information about the specified application. If information about an application other
10 than the one containing the requesting process is desired, then the attribute array must
11 contain a **PMIX_APPNUM** attribute identifying the desired target. Similarly, if information is
12 requested about an application in a job or session other than the one containing the requesting
13 process, then attributes identifying the target job and/or session must be provided.

14 **PMIX_NODE_INFO** "pmix.node.info" (bool)
15 Return information about the specified node. If information about a node other than the one
16 containing the requesting process is desired, then the attribute array must contain either the
17 **PMIX_NODEID** or **PMIX_HOSTNAME** attribute identifying the desired target.

18 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)
19 Process identifier Specifies the process ID whose information is being requested - e.g., a
20 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
21 request is for information on a specific process.

22 **PMIX_NAMESPACE** "pmix.namespace" (char*)
23 Namespace of the job. Specifies the namespace of the process whose information is being
24 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
25 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
26 information on a specific process.

27 **PMIX_RANK** "pmix.rank" (pmix_rank_t)
28 Process rank within the job. Specifies the rank of the process whose information is being
29 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
30 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
31 information on a specific process.

32 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)
33 Query list of supported attributes for specified APIs

34 **PMIX_CLIENT_ATTRIBUTES** "pmix.client.attrs" (bool)
35 Request attributes supported by the PMIx client library

36 **PMIX_SERVER_ATTRIBUTES** "pmix.srvr.attrs" (bool)
37 Request attributes supported by the PMIx server library

38 **PMIX_HOST_ATTRIBUTES** "pmix.host.attrs" (bool)
39 Request attributes supported by the host environment

1 **PMIX_TOOL_ATTRIBUTES** "pmix.setup.env" (bool)

2 Request attributes supported by the PMIx tool library functions

3 Note that inclusion of the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
4 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
5 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on
6 multiple specific processes therefore requires submitting multiple **pmix_query_t** structures,
7 each referencing one process.

8 PMIx libraries are not required to directly support any other attributes for this function. However,
9 any provided attributes must be passed to the host SMS daemon for processing, and the PMIx
10 library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client
11 process making the request.

13 Host environments that support this operation are required to support the following attributes as
14 qualifiers to the request:

15 **PMIX_PROCID** "pmix.procid" (**pmix_proc_t**)

16 Process identifier Specifies the process ID whose information is being requested - e.g., a
17 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
18 request is for information on a specific process.

19 **PMIX_NAMESPACE** "pmix.namespace" (**char***)

20 Namespace of the job. Specifies the namespace of the process whose information is being
21 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
22 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
23 information on a specific process.

24 **PMIX_RANK** "pmix.rank" (**pmix_rank_t**)

25 Process rank within the job. Specifies the rank of the process whose information is being
26 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
27 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
28 information on a specific process.

29 Note that inclusion of the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
30 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
31 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on
32 multiple specific processes therefore requires submitting multiple **pmix_query_t** structures,
33 each referencing one process.

Optional Attributes

34 The following attributes are optional for host environments that support this operation:

35 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (**char***)

36 Request a comma-delimited list of active namespaces.

1 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)
 2 Status of a specified, currently executing job.

3 **PMIX_QUERY_QUEUE_LIST** "pmix.qry.qlst" (char*)
 4 Request a comma-delimited list of scheduler queues.

5 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (TBD)
 6 Status of a specified scheduler queue.

7 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)
 8 Input namespace of the job whose information is being requested returns (
 9 pmix_data_array_t) an array of pmix_proc_info_t .

10 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lhtable" (char*)
 11 Input namespace of the job whose information is being requested returns (
 12 pmix_data_array_t) an array of pmix_proc_info_t for processes in job on same
 13 node.

14 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)
 15 Return a comma-delimited list of supported spawn attributes.

16 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)
 17 Return a comma-delimited list of supported debug attributes.

18 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)
 19 Return information on memory usage for the processes indicated in the qualifiers.

20 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
 21 Report only average values for sampled information.

22 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
 23 Report minimum and maximum values.

24 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 25 String identifier of the allocation whose status is being requested.

26 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
 27 Query number of seconds (uint32_t) remaining in allocation for the specified namespace.
 28

29 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)
 30 URI of the PMIx server to be contacted. Requests the URI of the specified PMIx server's
 31 PMIx connection. Defaults to requesting the information for the local PMIx server.

32 **PMIX_PROC_URI** "pmix.puri" (char*)
 33 URI containing contact information for a given process. Requests the URI of the specified
 34 PMIx server's out-of-band connection. Defaults to requesting the information for the local
 35 PMIx server.

▲-----▲

Description

Non-blocking form of the `PMIx_Query_info` API

7.1.4.1 Using `PMIx_Get` vs `PMIx_Query_info`

Both `PMIx_Get` and `PMIx_Query_info` can be used to retrieve information about the system. In general, the *get* operation should be used to retrieve:

- information provided by the host environment at time of job start. This includes information on the number of processes in the job, their location, and possibly their communication endpoints
- information posted by processes via the `PMIx_Put` function

This information is largely considered to be *static*, although this will not necessarily be true for environments supporting dynamic programming models or fault tolerance. Note that the `PMIx_Get` function only accesses information about execution environments - i.e., its scope is limited to values pertaining to a specific `session`, `job`, `application`, process, or node. It cannot be used to obtain information about areas such as the status of queues in the WLM.

In contrast, the *query* option should be used to access:

- system-level information (such as the available WLM queues) that would generally not be included in job-level information provided at job start
- dynamic information such as application and queue status, and resource utilization statistics. Note that the `PMIX_QUERY_REFRESH_CACHE` attribute must be provided on each query to ensure current data is returned
- information created post job start, such as process tables
- information requiring more complex search criteria than supported by the simpler `PMIx_Get` API
- queries focused on retrieving multi-attribute blocks of data with a single request, thus bypassing the single-key limitation of the `PMIx_Get` API

In theory, all information can be accessed via `PMIx_Query_info` as the local cache is typically the same datastore searched by `PMIx_Get`. However, in practice, the overhead associated with the *query* operation may (depending upon implementation) be higher than the simpler *get* operation due to the need to construct and process the more complex `pmix_query_t` structure. Thus, requests for a single key value are likely to be accomplished faster with `PMIx_Get` versus the *query* operation.

7.1.4.2 Accessing attribute support information

Information as to attributes supported by either the PMIx implementation or its host environment can be obtained via the `PMIx_Query_info_nb` API. The `PMIX_QUERY_ATTRIBUTE_SUPPORT` attribute must be listed as the first entry in the *keys* field of the `pmix_query_t` structure, followed by the name of the function whose attribute support is being requested - support for multiple functions can be requested simultaneously by simply adding

the function names to the array of *keys*. Function names *must* be given as user-level API names - e.g., “PMIx_Get”, “PMIx_server_setup_application”, or “PMIx_tool_connect_to_server”.

The desired levels (see 3.4.33) of attribute support are provided as qualifiers. Multiple levels can be requested simultaneously by simply adding elements to the *qualifiers* array. Each qualifier should contain the desired level attribute with the boolean value set to indicate whether or not that level is to be included in the returned information. Failure to provide any levels is equivalent to a request for all levels.

Unlike other queries, queries for attribute support can result in the number of returned `pmix_info_t` structures being different from the number of queries. Each element in the returned array will correspond to a pair of specified attribute level and function in the query, where the *key* is the function and the *value* contains a `pmix_data_array_t` of `pmix_info_t`. Each element of the array is marked by a *key* indicating the requested attribute *level* with a *value* composed of a `pmix_data_array_t` of `pmix_regattr_t`, each describing a supported attribute for that function, as illustrated in Fig. 7.1 below where the requestor asked for supported attributes of `PMIx_Get` at the *client* and *server* levels, plus attributes of `PMIx_Allocation_request` at all levels:




Figure 7.1.: Returned information hierarchy for attribute support request

The array of returned structures, and their child arrays, are subject to the return rules for the `PMIx_Query_info_nb` API. For example, a request for supported attributes of the `PMIx_Get` function that includes the *host* level will return values for the *client* and *server* levels, plus an array element with a *key* of `PMIX_HOST_ATTRIBUTES` and a value type of `PMIX_UNDEF` indicating that no attributes are supported at that level.

7.2 Allocation Requests

This section defines functionality to request new allocations from the RM, and request modifications to existing allocations. These are primarily used in the following scenarios:

- 1 • *Evolving* applications that dynamically request and return resources as they execute
- 2 • *Malleable* environments where the scheduler redirects resources away from executing
- 3 applications for higher priority jobs or load balancing
- 4 • *Resilient* applications that need to request replacement resources in the face of failures
- 5 • *Rigid* jobs where the user has requested a static allocation of resources for a fixed period of time,
- 6 but realizes that they underestimated their required time while executing
- 7 PMIx attempts to address this range of use-cases with a flexible API.

8 7.2.1 PMIx_Allocation_request

9 Summary

10 Request an allocation operation from the host resource manager.

11 Format

PMIx v3.0

```

12 pmix_status_t
13 PMIx_Allocation_request(pmix_alloc_directive_t directive,
14                        pmix_info_t info[], size_t ninfo);

```

- 15 **IN directive**
Allocation directive (handle)
- 16 **IN info**
Array of `pmix_info_t` structures (array of handles)
- 17 **IN ninfo**
Number of elements in the *info* array (integer)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS**, indicating that the request was processed and returned *success*
- 23 • a PMIx error constant indicating either an error in the input or that the request was refused

Required Attributes

24 PMIx libraries are not required to directly support any attributes for this function. However, any
 25 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
 26 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process making
 27 the request.

29 Host environments that implement support for this operation are required to support the following
 30 attributes:

31 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)

1 Provide a string identifier for this allocation request which can later be used to query status
2 of the request.

3 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)
4 The number of nodes.

5 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
6 Number of cpus.

7 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
8 Time in seconds.



▼----- Optional Attributes -----▼

9 The following attributes are optional for host environments that support this operation:

10 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
11 Regular expression of the specific nodes.

12 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
13 Regular expression of the number of cpus for each node.

14 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
15 Regular expression of the specific cpus indicating the cpus involved.

16 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
17 Number of Megabytes.

18 **PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)
19 Array of **pmix_info_t** describing requested network resources. This must include at
20 least: **PMIX_ALLOC_NETWORK_ID**, **PMIX_ALLOC_NETWORK_TYPE**, and
21 **PMIX_ALLOC_NETWORK_ENDPTS**, plus whatever other descriptors are desired.

22 **PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)
23 The key to be used when accessing this requested network allocation. The allocation will be
24 returned/stored as a **pmix_data_array_t** of **pmix_info_t** indexed by this key and
25 containing at least one entry with the same key and the allocated resource description. The
26 type of the included value depends upon the network support. For example, a TCP allocation
27 might consist of a comma-delimited string of socket ranges such as
28 "32000-32100,33005,38123-38146". Additional entries will consist of any provided
29 resource request directives, along with their assigned values. Examples include:
30 **PMIX_ALLOC_NETWORK_TYPE** - the type of resources provided;
31 **PMIX_ALLOC_NETWORK_PLANE** - if applicable, what plane the resources were assigned
32 from; **PMIX_ALLOC_NETWORK_QOS** - the assigned QoS; **PMIX_ALLOC_BANDWIDTH** -
33 the allocated bandwidth; **PMIX_ALLOC_NETWORK_SEC_KEY** - a security key for the
34 requested network allocation. NOTE: the assigned values may differ from those requested,
35 especially if **PMIX_INFO_REQD** was not set in the request.

36 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)

1 Mbits/sec.

2 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)

3 Quality of service level.

4 **PMIX_ALLOC_NETWORK_TYPE** "pmix.alloc.nettype" (char*)

5 Type of desired transport (e.g., "tcp", "udp")

6 **PMIX_ALLOC_NETWORK_PLANE** "pmix.alloc.netplane" (char*)

7 ID string for the NIC (aka *plane*) to be used for this allocation (e.g., CIDR for Ethernet)

8 **PMIX_ALLOC_NETWORK_ENDPTS** "pmix.alloc.endpts" (size_t)

9 Number of endpoints to allocate per process

10 **PMIX_ALLOC_NETWORK_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)

11 Number of endpoints to allocate per node

12 **PMIX_ALLOC_NETWORK_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)

13 Network security key



14 **Description**

15 Request an allocation operation from the host resource manager. Several broad categories are
16 envisioned, including the ability to:

- 17 • Request allocation of additional resources, including memory, bandwidth, and compute. This
18 should be accomplished in a non-blocking manner so that the application can continue to
19 progress while waiting for resources to become available. Note that the new allocation will be
20 disjoint from (i.e., not affiliated with) the allocation of the requestor - thus the termination of one
21 allocation will not impact the other.
- 22 • Extend the reservation on currently allocated resources, subject to scheduling availability and
23 priorities. This includes extending the time limit on current resources, and/or requesting
24 additional resources be allocated to the requesting job. Any additional allocated resources will be
25 considered as part of the current allocation, and thus will be released at the same time.
- 26 • Return no-longer-required resources to the scheduler. This includes the "loan" of resources back
27 to the scheduler with a promise to return them upon subsequent request.

28 **7.2.2 PMIx_Allocation_request_nb**

29 **Summary**

30 Request an allocation operation from the host resource manager.

1
PMIx v2.0

Format

C

```

2 pmix_status_t
3 PMIx_Allocation_request_nb(pmix_alloc_directive_t directive,
4                             pmix_info_t info[], size_t ninfo,
5                             pmix_info_cbfunc_t cbfunc, void *cbdata);

```

C

- 6 **IN directive**
Allocation directive (handle)
- 7
- 8 **IN info**
Array of `pmix_info_t` structures (array of handles)
- 9
- 10 **IN ninfo**
Number of elements in the *info* array (integer)
- 11
- 12 **IN cbfunc**
Callback function `pmix_info_cbfunc_t` (function reference)
- 13
- 14 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 15

16 Returns one of the following:

- 17 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
18 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
19 function prior to returning from the API.
- 20 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
21 returned *success* - the *cbfunc* will *not* be called
- 22 • a PMIx error constant indicating either an error in the input or that the request was immediately
23 processed and failed - the *cbfunc* will *not* be called

Required Attributes

24 PMIx libraries are not required to directly support any attributes for this function. However, any
25 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
26 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process making
27 the request.


28

29 Host environments that implement support for this operation are required to support the following
30 attributes:

31 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)
32 Provide a string identifier for this allocation request which can later be used to query status
33 of the request.

34 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)

1 The number of nodes.
2 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
3 Number of cpus.
4 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
5 Time in seconds.



6 The following attributes are optional for host environments that support this operation:

7 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
8 Regular expression of the specific nodes.

9 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
10 Regular expression of the number of cpus for each node.

11 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
12 Regular expression of the specific cpus indicating the cpus involved.

13 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
14 Number of Megabytes.

15 **PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)
16 Array of `pmix_info_t` describing requested network resources. This must include at
17 least: `PMIX_ALLOC_NETWORK_ID`, `PMIX_ALLOC_NETWORK_TYPE`, and
18 **PMIX_ALLOC_NETWORK_ENDPTS**, plus whatever other descriptors are desired.

19 **PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)
20 The key to be used when accessing this requested network allocation. The allocation will be
21 returned/stored as a `pmix_data_array_t` of `pmix_info_t` indexed by this key and
22 containing at least one entry with the same key and the allocated resource description. The
23 type of the included value depends upon the network support. For example, a TCP allocation
24 might consist of a comma-delimited string of socket ranges such as
25 "32000-32100,33005,38123-38146". Additional entries will consist of any provided
26 resource request directives, along with their assigned values. Examples include:
27 **PMIX_ALLOC_NETWORK_TYPE** - the type of resources provided;
28 **PMIX_ALLOC_NETWORK_PLANE** - if applicable, what plane the resources were assigned
29 from; **PMIX_ALLOC_NETWORK_QOS** - the assigned QoS; **PMIX_ALLOC_BANDWIDTH** -
30 the allocated bandwidth; **PMIX_ALLOC_NETWORK_SEC_KEY** - a security key for the
31 requested network allocation. NOTE: the assigned values may differ from those requested,
32 especially if **PMIX_INFO_REQD** was not set in the request.

33 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)
34 Mbits/sec.

35 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)

1 Quality of service level.

2 **PMIX_ALLOC_NETWORK_TYPE** "pmix.alloc.nettype" (char*)

3 Type of desired transport (e.g., "tcp", "udp")

4 **PMIX_ALLOC_NETWORK_PLANE** "pmix.alloc.netplane" (char*)

5 ID string for the NIC (aka *plane*) to be used for this allocation (e.g., CIDR for Ethernet)

6 **PMIX_ALLOC_NETWORK_ENDPTS** "pmix.alloc.endpts" (size_t)

7 Number of endpoints to allocate per process

8 **PMIX_ALLOC_NETWORK_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)

9 Number of endpoints to allocate per node

10 **PMIX_ALLOC_NETWORK_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)

11 Network security key



12 **Description**

13 Non-blocking form of the [PMIx_Allocation_request](#) API.

14 **7.3 Job Control**

15 This section defines APIs that enable the application and host environment to coordinate the
16 response to failures and other events. This can include requesting termination of the entire job or a
17 subset of processes within a job, but can also be used in combination with other PMIx capabilities
18 (e.g., allocation support and event notification) for more nuanced responses. For example, an
19 application notified of an incipient over-temperature condition on a node could use the
20 [PMIx_Allocation_request_nb](#) interface to request replacement nodes while
21 simultaneously using the [PMIx_Job_control_nb](#) interface to direct that a checkpoint event be
22 delivered to all processes in the application. If replacement resources are not available, the
23 application might use the [PMIx_Job_control_nb](#) interface to request that the job continue at
24 a lower power setting, perhaps sufficient to avoid the over-temperature failure.

25 The job control APIs can also be used by an application to register itself as available for preemption
26 when operating in an environment such as a cloud or where incentives, financial or otherwise, are
27 provided to jobs willing to be preempted. Registration can include attributes indicating how many
28 resources are being offered for preemption (e.g., all or only some portion), whether the application
29 will require time to prepare for preemption, etc. Jobs that request a warning will receive an event
30 notifying them of an impending preemption (possibly including information as to the resources that
31 will be taken away, how much time the application will be given prior to being preempted, whether
32 the preemption will be a suspension or full termination, etc.) so they have an opportunity to save
33 their work. Once the application is ready, it calls the provided event completion callback function to
34 indicate that the SMS is free to suspend or terminate it, and can include directives regarding any
35 desired restart.

1 7.3.1 PMIx_Job_control

2 Summary

3 Request a job control action.

4 Format

PMIx v3.0

```
5 pmix_status_t
6 PMIx_Job_control(const pmix_proc_t targets[], size_t ntargets,
7                  const pmix_info_t directives[], size_t ndirs)
```

8 IN targets

9 Array of proc structures (array of handles)

10 IN ntargets

11 Number of element in the *targets* array (integer)

12 IN directives

13 Array of info structures (array of handles)

14 IN ndirs

15 Number of element in the *directives* array (integer)

16 IN cbfunc

17 Callback function [pmix_info_cbfunc_t](#) (function reference)

18 IN cbdata

19 Data to be passed to the callback function (memory reference)

20 Returns one of the following:

- 21 • **PMIX_SUCCESS**, indicating that the request was processed by the host environment and
22 returned *success*
- 23 • a PMIx error constant indicating either an error in the input or that the request was refused

Required Attributes

24 PMIx libraries are not required to directly support any attributes for this function. However, any
25 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
26 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process making
27 the request.

29 Host environments that implement support for this operation are required to support the following
30 attributes:

31 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)

32 Provide a string identifier for this request. The user can provide an identifier for the
33 requested operation, thus allowing them to later request status of the operation or to
34 terminate it. The host, therefore, shall track it with the request for future reference.

1 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)
 2 Pause the specified processes.

3 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)
 4 Resume ("un-pause") the specified processes.

5 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
 6 Forcibly terminate the specified processes and cleanup.

7 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
 8 Send given signal to specified processes.

9 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
 10 Politely terminate the specified processes.

11 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)
 12 Comma-delimited list of files to be removed upon process termination


13 **PMIX_REGISTER_CLEANUP_DIR** "pmix.reg.cleanupdir" (char*)
 14 Comma-delimited list of directories to be removed upon process termination

15 **PMIX_CLEANUP_RECURSIVE** "pmix.clnup.recurse" (bool)
 16 Recursively cleanup all subdirectories under the specified one(s)

17 **PMIX_CLEANUP_EMPTY** "pmix.clnup.empty" (bool)
 18 Only remove empty subdirectories

19 **PMIX_CLEANUP_IGNORE** "pmix.clnup.ignore" (char*)
 20 Comma-delimited list of filenames that are not to be removed

21 **PMIX_CLEANUP_LEAVE_TOPDIR** "pmix.clnup.lvtop" (bool)
 22 When recursively cleaning subdirectories, do not remove the top-level directory (the one
 23 given in the cleanup request)



24 The following attributes are optional for host environments that support this operation:

25 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
 26 Cancel the specified request - the provided request ID must match the
 27 **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control** . An ID of
 28 **NULL** implies cancel all requests from this requestor.

29 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
 30 Restart the specified processes using the given checkpoint ID.

31 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
 32 Checkpoint the specified processes and assign the given ID to it.

33 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)

1 Use event notification to trigger a process checkpoint.

2 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)

3 Use the given signal to trigger a process checkpoint.

4 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)

5 Time in seconds to wait for a checkpoint to complete.

6 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**

7 "pmix.jctrl.ckmethod" (pmix_data_array_t)

8 Array of **pmix_info_t** declaring each method and value supported by this application.

9 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)

10 Regular expression identifying nodes that are to be provisioned.

11 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)

12 Name of the image that is to be provisioned.

13 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)

14 Indicate that the job can be pre-empted.



Description

Request a job control action. The *targets* array identifies the processes to which the requested job control action is to be applied. A **NULL** value can be used to indicate all processes in the caller's namespace. The use of **PMIX_RANK_WILDCARD** can also be used to indicate that all processes in the given namespace are to be included.

The directives are provided as **pmix_info_t** structures in the *directives* array. The callback function provides a *status* to indicate whether or not the request was granted, and to provide some information as to the reason for any denial in the **pmix_info_cbfnc_t** array of **pmix_info_t** structures.

24 7.3.2 PMIx_Job_control_nb

25 Summary

26 Request a job control action.

1
PMIx v2.0

Format

C

```

2 pmix_status_t
3 PMIx_Job_control_nb(const pmix_proc_t targets[], size_t ntargets,
4                     const pmix_info_t directives[], size_t ndirs,
5                     pmix_info_cbfunc_t cbfunc, void *cbdata)

```

C

- 6 **IN targets**
Array of proc structures (array of handles)
- 7
- 8 **IN ntargets**
Number of element in the *targets* array (integer)
- 9
- 10 **IN directives**
Array of info structures (array of handles)
- 11
- 12 **IN ndirs**
Number of element in the *directives* array (integer)
- 13
- 14 **IN cbfunc**
Callback function `pmix_info_cbfunc_t` (function reference)
- 15
- 16 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 17

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
21 function prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will *not* be called
- 24 • a PMIx error constant indicating either an error in the input or that the request was immediately
25 processed and failed - the *cbfunc* will *not* be called

Required Attributes

26 PMIx libraries are not required to directly support any attributes for this function. However, any
27 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
28 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process making
29 the request.

30
31 Host environments that implement support for this operation are required to support the following
32 attributes:

33 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)

1 Provide a string identifier for this request. The user can provide an identifier for the
2 requested operation, thus allowing them to later request status of the operation or to
3 terminate it. The host, therefore, shall track it with the request for future reference.

4 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)

5 Pause the specified processes.

6 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)

7 Resume ("un-pause") the specified processes.

8 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)

9 Forcibly terminate the specified processes and cleanup.

10 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)

11 Send given signal to specified processes.

12 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)

13 Politely terminate the specified processes.

14 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)

15 Comma-delimited list of files to be removed upon process termination

16 **PMIX_REGISTER_CLEANUP_DIR** "pmix.reg.cleanupdir" (char*)

17 Comma-delimited list of directories to be removed upon process termination

18 **PMIX_CLEANUP_RECURSIVE** "pmix.clnup.recurse" (bool)

19 Recursively cleanup all subdirectories under the specified one(s)

20 **PMIX_CLEANUP_EMPTY** "pmix.clnup.empty" (bool)


21 Only remove empty subdirectories

22 **PMIX_CLEANUP_IGNORE** "pmix.clnup.ignore" (char*)

23 Comma-delimited list of filenames that are not to be removed

24 **PMIX_CLEANUP_LEAVE_TOPDIR** "pmix.clnup.lvtop" (bool)

25 When recursively cleaning subdirectories, do not remove the top-level directory (the one
26 given in the cleanup request)



Optional Attributes

27 The following attributes are optional for host environments that support this operation:

28 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)

29 Cancel the specified request - the provided request ID must match the

30 **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control** . An ID of

31 **NULL** implies cancel all requests from this requestor.

32 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)

33 Restart the specified processes using the given checkpoint ID.

1 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
 2 Checkpoint the specified processes and assign the given ID to it.

3 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
 4 Use event notification to trigger a process checkpoint.

5 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
 6 Use the given signal to trigger a process checkpoint.

7 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
 8 Time in seconds to wait for a checkpoint to complete.

9 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
 10 "pmix.jctrl.ckmethod" (pmix_data_array_t)
 11 Array of **pmix_info_t** declaring each method and value supported by this application.

12 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
 13 Regular expression identifying nodes that are to be provisioned.

14 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
 15 Name of the image that is to be provisioned.

16 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
 17 Indicate that the job can be pre-empted.



18 Description

19 Non-blocking form of the **PMIx_Job_control** API. The *targets* array identifies the processes
 20 to which the requested job control action is to be applied. A **NULL** value can be used to indicate all
 21 processes in the caller's namespace. The use of **PMIX_RANK_WILDCARD** can also be used to
 22 indicate that all processes in the given namespace are to be included.

23 The directives are provided as **pmix_info_t** structures in the *directives* array. The callback
 24 function provides a *status* to indicate whether or not the request was granted, and to provide some
 25 information as to the reason for any denial in the **pmix_info_cbfunc_t** array of
 26 **pmix_info_t** structures.

27 7.4 Process and Job Monitoring

28 In addition to external faults, a common problem encountered in HPC applications is a failure to
 29 make progress due to some internal conflict in the computation. These situations can result in a
 30 significant waste of resources as the SMS is unaware of the problem, and thus cannot terminate the
 31 job. Various watchdog methods have been developed for detecting this situation, including
 32 requiring a periodic "heartbeat" from the application and monitoring a specified file for changes in
 33 size and/or modification time.

1 At the request of SMS vendors and members, a monitoring support interface has been included in
2 the PMIx v2 standard. The defined API allows applications to request monitoring, directing what is
3 to be monitored, the frequency of the associated check, whether or not the application is to be
4 notified (via the event notification subsystem) of stall detection, and other characteristics of the
5 operation. In addition, heartbeat and file monitoring methods have been included in the PRI but are
6 active only when requested.

7 7.4.1 PMIx_Process_monitor

8 Summary

9 Request that application processes be monitored.

10 Format

PMIx v3.0

```
11 pmix_status_t  
12 PMIx_Process_monitor(const pmix_info_t *monitor, pmix_status_t error,  
13                      const pmix_info_t directives[], size_t ndirs)
```

14 **IN monitor**
15 info (handle)
16 **IN error**
17 status (integer)
18 **IN directives**
19 Array of info structures (array of handles)
20 **IN ndirs**
21 Number of elements in the *directives* array (integer)

22 Returns one of the following:

- 23 • **PMIX_SUCCESS**, indicating that the request was processed and returned *success*
- 24 • a PMIx error constant indicating either an error in the input or that the request was refused

Optional Attributes

25 The following attributes may be implemented by a PMIx library or by the host environment. If
26 supported by the PMIx server library, then the library must not pass the supported attributes to the
27 host environment. All attributes not directly supported by the server library must be passed to the
28 host environment if it supports this operation, and the library is *required* to add the
29 **PMIX_USERID** and the **PMIX_GRPID** attributes of the requesting process:

30 **PMIX_MONITOR_ID** "pmix.monitor.id" (char*)

31 Provide a string identifier for this request.

32 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)

33 Identifier to be canceled (**NULL** means cancel all monitoring for this process).

1 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)
 2 The application desires to control the response to a monitoring event.

3 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)
 4 Register to have the PMIx server monitor the requestor for heartbeats.

5 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)
 6 Time in seconds before declaring heartbeat missed.

7 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrop" (uint32_t)
 8 Number of heartbeats that can be missed before generating the event.

9 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)
 10 Register to monitor file for signs of life.

11 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)
 12 Monitor size of given file is growing to determine if the application is running.

13 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)
 14 Monitor time since last access of given file to determine if the application is running.

15 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
 16 Monitor time since last modified of given file to determine if the application is running.

17 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
 18 Time in seconds between checking the file.

19 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)
 20 Number of file checks that can be missed before generating the event.



21 **Description**

22 Request that application processes be monitored via several possible methods. For example, that
 23 the server monitor this process for periodic heartbeats as an indication that the process has not
 24 become “wedged”. When a monitor detects the specified alarm condition, it will generate an event
 25 notification using the provided error code and passing along any available relevant information. It
 26 is up to the caller to register a corresponding event handler.

27 The *monitor* argument is an attribute indicating the type of monitor being requested. For example,
 28 **PMIX_MONITOR_FILE** to indicate that the requestor is asking that a file be monitored.

29 The *error* argument is the status code to be used when generating an event notification alerting that
 30 the monitor has been triggered. The range of the notification defaults to
 31 **PMIX_RANGE_NAMESPACE**. This can be changed by providing a **PMIX_RANGE** directive.

32 The *directives* argument characterizes the monitoring request (e.g., monitor file size) and frequency
 33 of checking to be done

1 7.4.2 PMIx_Process_monitor_nb

2 Summary

3 Request that application processes be monitored.

4 Format

PMIx v2.0

C

5 `pmix_status_t`

```
6 PMIx_Process_monitor_nb(const pmix_info_t *monitor, pmix_status_t error,  
7                         const pmix_info_t directives[], size_t ndirs,  
8                         pmix_info_cbfunc_t cbfunc, void *cbdata)
```

C

9 **IN monitor**

10 info (handle)

11 **IN error**

12 status (integer)

13 **IN directives**

14 Array of info structures (array of handles)

15 **IN ndirs**

16 Number of elements in the *directives* array (integer)

17 **IN cbfunc**

18 Callback function `pmix_info_cbfunc_t` (function reference)

19 **IN cbdata**

20 Data to be passed to the callback function (memory reference)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
23 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
24 function prior to returning from the API.
- 25 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
26 returned *success* - the *cbfunc* will *not* be called
- 27 • a PMIx error constant indicating either an error in the input or that the request was immediately
28 processed and failed - the *cbfunc* will *not* be called

Optional Attributes

29 The following attributes may be implemented by a PMIx library or by the host environment. If
30 supported by the PMIx server library, then the library must not pass the supported attributes to the
31 host environment. All attributes not directly supported by the server library must be passed to the
32 host environment if it supports this operation, and the library is *required* to add the
33 **PMIX_USERID** and the **PMIX_GRPID** attributes of the requesting process:

34 **PMIX_MONITOR_ID** "pmix.monitor.id" (char*)

1 Provide a string identifier for this request.

2 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)

3 Identifier to be canceled (NULL means cancel all monitoring for this process).

4 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)

5 The application desires to control the response to a monitoring event.

6 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)

7 Register to have the PMIx server monitor the requestor for heartbeats.

8 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)

9 Time in seconds before declaring heartbeat missed.

10 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrop" (uint32_t)

11 Number of heartbeats that can be missed before generating the event.

12 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)

13 Register to monitor file for signs of life.

14 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)

15 Monitor size of given file is growing to determine if the application is running.

16 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)

17 Monitor time since last access of given file to determine if the application is running.

18 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)

19 Monitor time since last modified of given file to determine if the application is running.

20 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)

21 Time in seconds between checking the file.

22 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)

23 Number of file checks that can be missed before generating the event.




24 Description

25 Non-blocking form of the [PMIx_Process_monitor](#) API. The *cbfunc* function provides a
26 *status* to indicate whether or not the request was granted, and to provide some information as to the
27 reason for any denial in the [pmix_info_cbfunc_t](#) array of [pmix_info_t](#) structures.

28 7.4.3 PMIx_Heartbeat

29 Summary

30 Send a heartbeat to the PMIx server library

1 **Format**
PMIx v2.0  

2 **PMIx_Heartbeat (void)**
  



3 **Description**
 4 A simplified macro wrapping **PMIx_Process_monitor_nb** that sends a heartbeat to the
 5 PMIx server library.




6 7.5 Logging

7 The logging interface supports posting information by applications and SMS elements to persistent
 8 storage. This function is *not* intended for output of computational results, but rather for reporting
 9 status and saving state information such as inserting computation progress reports into the
 10 application's SMS job log or error reports to the local syslog.

11 7.5.1 PMIx_Log

12 **Summary**
 13 Log data to a data service.

14 **Format**
PMIx v3.0  

15 **pmix_status_t**
 16 **PMIx_Log**(const **pmix_info_t** data[], **size_t** ndata,
 17 **const pmix_info_t** directives[], **size_t** ndirs)
  

18 **IN data**
 19 Array of info structures (array of handles)
 20 **IN ndata**
 21 Number of elements in the *data* array (**size_t**)
 22 **IN directives**
 23 Array of info structures (array of handles)
 24 **IN ndirs**
 25 Number of elements in the *directives* array (**size_t**)

26 Return codes are one of the following:
 27 **PMIX_SUCCESS** The logging request was successful.
 28 **PMIX_ERR_BAD_PARAM** The logging request contains at least one incorrect entry.
 29 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation or host environment does not
 30 support this function.

Required Attributes

1 If the PMIx library does not itself perform this operation, then it is required to pass any attributes
2 provided by the client to the host environment for processing. In addition, it must include the
3 following attributes in the passed *info* array:

4 **PMIX_USERID** "pmix.euid" (uint32_t)
5 Effective user id.

6 **PMIX_GRPID** "pmix.egid" (uint32_t)
7 Effective group id.

8
9 Host environments or PMIx libraries that implement support for this operation are required to
10 support the following attributes:

11 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)
12 Log string to **stderr**.

13 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)
14 Log string to **stdout**.

15 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)
16 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
17 otherwise to local syslog

18 **PMIX_LOG_LOCAL_SYSLOG** "pmix.log.lsys" (char*)
19 Log data to local syslog. Defaults to **ERROR** priority.

20 **PMIX_LOG_GLOBAL_SYSLOG** "pmix.log.gsys" (char*)
21 Forward data to system "gateway" and log msg to that syslog Defaults to **ERROR** priority.

22 **PMIX_LOG_SYSLOG_PRI** "pmix.log.syspri" (int)
23 Syslog priority level

24 **PMIX_LOG_ONCE** "pmix.log.once" (bool)
25 Only log this once with whichever channel can first support it, taking the channels in priority
26 order

Optional Attributes

27 The following attributes are optional for host environments or PMIx libraries that support this
28 operation:

29 **PMIX_LOG_SOURCE** "pmix.log.source" (pmix_proc_t*)
30 ID of source of the log request

31 **PMIX_LOG_TIMESTAMP** "pmix.log.tstamp" (time_t)
32 Timestamp for log report

1 **PMIX_LOG_GENERATE_TIMESTAMP** "pmix.log.gtstmp" (bool)
2 Generate timestamp for log

3 **PMIX_LOG_TAG_OUTPUT** "pmix.log.tag" (bool)
4 Label the output stream with the channel name (e.g., "stdout")

5 **PMIX_LOG_TIMESTAMP_OUTPUT** "pmix.log.tsout" (bool)
6 Print timestamp in output string

7 **PMIX_LOG_XML_OUTPUT** "pmix.log.xml" (bool)
8 Print the output stream in XML format

9 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)
10 Log via email based on **pmix_info_t** containing directives.

11 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)
12 Comma-delimited list of email addresses that are to receive the message.

13 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)
14 Subject line for email.

15 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)
16 Message to be included in email.

17 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)
18 Log the provided information to the host environment's job record

19 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)
20 Store the log data in a global data store (e.g., database)



21 **Description**

22 Log data subject to the services offered by the host environment. The data to be logged is provided
23 in the *data* array. The (optional) *directives* can be used to direct the choice of logging channel.

24 **Advice to users**

25 It is strongly recommended that the **PMIx_Log** API not be used by applications for streaming data
26 as it is not a "performant" transport and can perturb the application since it involves the local PMIx
27 server and host SMS daemon. Note that a return of **PMIX_SUCCESS** only denotes that the data
28 was successfully handed to the appropriate system call (for local channels) or the host environment
and does not indicate receipt at the final destination.

29 **7.5.2 PMIx_Log_nb**

30 **Summary**

31 Log data to a data service.

1
PMIx v2.0

Format

C

```

2 pmix_status_t
3 PMIx_Log_nb(const pmix_info_t data[], size_t ndata,
4             const pmix_info_t directives[], size_t ndirs,
5             pmix_op_cbfunc_t cbfunc, void *cbdata)

```

C

- 6 **IN data**
Array of info structures (array of handles)
- 7
- 8 **IN ndata**
Number of elements in the *data* array (**size_t**)
- 9
- 10 **IN directives**
Array of info structures (array of handles)
- 11
- 12 **IN ndirs**
Number of elements in the *directives* array (**size_t**)
- 13
- 14 **IN cbfunc**
Callback function **pmix_op_cbfunc_t** (function reference)
- 15
- 16 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 17

18 Return codes are one of the following:

- 19 **PMIX_SUCCESS** The logging request is valid and is being processed. The resulting status from
- 20 the operation will be provided in the callback function. Note that the library must not invoke
- 21 the callback function prior to returning from the API.
- 22 **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
- 23 returned *success* - the *cbfunc* will *not* be called
- 24 **PMIX_ERR_BAD_PARAM** The logging request contains at least one incorrect entry that prevents
- 25 it from being processed. The callback function will not be called.
- 26 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function. The
- 27 callback function will not be called.

Required Attributes

28 If the PMIx library does not itself perform this operation, then it is required to pass any attributes
29 provided by the client to the host environment for processing. In addition, it must include the
30 following attributes in the passed *info* array:

- 31 **PMIX_USERID** "pmix.euid" (**uint32_t**)
32 Effective user id.
- 33 **PMIX_GRPID** "pmix.egid" (**uint32_t**)
34 Effective group id.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33

Host environments or PMIx libraries that implement support for this operation are required to support the following attributes:

- PMIX_LOG_STDERR** "pmix.log.stderr" (char*)
Log string to **stderr**.
- PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)
Log string to **stdout**.
- PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)
Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available, otherwise to local syslog
- PMIX_LOG_LOCAL_SYSLOG** "pmix.log.lsys" (char*)
Log data to local syslog. Defaults to **ERROR** priority.
- PMIX_LOG_GLOBAL_SYSLOG** "pmix.log.gsys" (char*)
Forward data to system "gateway" and log msg to that syslog Defaults to **ERROR** priority.
- PMIX_LOG_SYSLOG_PRI** "pmix.log.syspri" (int)
Syslog priority level
- PMIX_LOG_ONCE** "pmix.log.once" (bool)
Only log this once with whichever channel can first support it, taking the channels in priority order

▲-----▲
▼-----▼ **Optional Attributes** -----▼

The following attributes are optional for host environments or PMIx libraries that support this operation:

- PMIX_LOG_SOURCE** "pmix.log.source" (pmix_proc_t*)
ID of source of the log request
- PMIX_LOG_TIMESTAMP** "pmix.log.tstamp" (time_t)
Timestamp for log report
- PMIX_LOG_GENERATE_TIMESTAMP** "pmix.log.gtstamp" (bool)
Generate timestamp for log
- PMIX_LOG_TAG_OUTPUT** "pmix.log.tag" (bool)
Label the output stream with the channel name (e.g., "stdout")
- PMIX_LOG_TIMESTAMP_OUTPUT** "pmix.log.tsout" (bool)
Print timestamp in output string
- PMIX_LOG_XML_OUTPUT** "pmix.log.xml" (bool)
Print the output stream in XML format

1 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)
2 Log via email based on **pmix_info_t** containing directives.
3 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)
4 Comma-delimited list of email addresses that are to receive the message.
5 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)
6 Subject line for email.
7 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)
8 Message to be included in email.
9 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)
10 Log the provided information to the host environment's job record
11 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)
12 Store the log data in a global data store (e.g., database)



Description

13 Log data subject to the services offered by the host environment. The data to be logged is provided
14 in the *data* array. The (optional) *directives* can be used to direct the choice of logging channel. The
15 callback function will be executed when the log operation has been completed. The *data* and
16 *directives* arrays must be maintained until the callback is provided.
17

Advice to users

18 It is strongly recommended that the **PMIx_Log_nb** API not be used by applications for streaming
19 data as it is not a “performant” transport and can perturb the application since it involves the local
20 PMIx server and host SMS daemon. Note that a return of **PMIX_SUCCESS** only denotes that the
21 data was successfully handed to the appropriate system call (for local channels) or the host
22 environment and does not indicate receipt at the final destination.



CHAPTER 8

Event Notification

1 This chapter defines the PMIx event notification system. These interfaces are designed to support
2 the reporting of events to/from clients and servers, and between library layers within a single
3 process.

4 8.1 Notification and Management

5 PMIx event notification provides an asynchronous out-of-band mechanism for communicating
6 events between application processes and/or elements of the SMS. Its uses span a wide range that
7 includes fault notification, coordination between multiple programming libraries within a single
8 process, and workflow orchestration for non-synchronous programming models. Events can be
9 divided into two distinct classes:

- 10 • *Job-specific events* directly relate to a job executing within the session, such as a debugger
11 attachment, process failure within a related job, or events generated by an application process.
12 Events in this category are to be immediately delivered to the PMIx server library for relay to the
13 related local processes.
- 14 • *Environment events* indirectly relate to a job but do not specifically target the job itself. This
15 category includes SMS-generated events such as Error Check and Correction (ECC) errors,
16 temperature excursions, and other non-job conditions that might directly affect a session's
17 resources, but would never include an event generated by an application process. Note that
18 although these do potentially impact the session's jobs, they are not directly tied to those jobs.
19 Thus, events in this category are to be delivered to the PMIx server library only upon request.

20 Both SMS elements and applications can register for events of either type.

▼ Advice to PMIx library implementers ▼

21 Race conditions can cause the registration to come after events of possible interest (e.g., a memory
22 ECC event that occurs after start of execution but prior to registration, or an application process
23 generating an event prior to another process registering to receive it). SMS vendors are *requested* to
24 cache environment events for some time to mitigate this situation, but are not *required* to do so.
25 However, PMIx implementers are *required* to cache all events received by the PMIx server library
26 and to deliver them to registering clients in the same order in which they were received

Advice to users

1 Applications must be aware that they may not receive environment events that occur prior to
2 registration, depending upon the capabilities of the host SMS.

3 The generator of an event can specify the *target range* for delivery of that event. Thus, the generator
4 can choose to limit notification to processes on the local node, processes within the same job as the
5 generator, processes within the same allocation, other threads within the same process, only the
6 SMS (i.e., not to any application processes), all application processes, or to a custom range based
7 on specific process identifiers. Only processes within the given range that register for the provided
8 event code will be notified. In addition, the generator can use attributes to direct that the event not
9 be delivered to any default event handlers, or to any multi-code handler (as defined below).

10 Event notifications provide the process identifier of the source of the event plus the event code and
11 any additional information provided by the generator. When an event notification is received by a
12 process, the registered handlers are scanned for their event code(s), with matching handlers
13 assembled into an *event chain* for servicing. Note that users can also specify a *source range* when
14 registering an event (using the same range designators described above) to further limit when they
15 are to be invoked. When assembled, PMIx event chains are ordered based on both the specificity of
16 the event handler and user directives at time of handler registration. By default, handlers are
17 grouped into three categories based on the number of event codes that can trigger the callback:

- 18 ● *single-code* handlers are serviced first as they are the most specific. These are handlers that are
19 registered against one specific event code.
- 20 ● *multi-code* handlers are serviced once all single-code handlers have completed. The handler will
21 be included in the chain upon receipt of an event matching any of the provided codes.
- 22 ● *default* handlers are serviced once all multi-code handlers have completed. These handlers are
23 always included in the chain unless the generator specifically excludes them.

24 Users can specify the callback order of a handler within its category at the time of registration.
25 Ordering can be specified either by providing the relevant returned event handler registration ID or
26 using event handler names, if the user specified an event handler name when registering the
27 corresponding event. Thus, users can specify that a given handler be executed before or after
28 another handler should both handlers appear in an event chain (the ordering is ignored if the other
29 handler isn't included). Note that ordering does not imply immediate relationships. For example,
30 multiple handlers registered to be serviced after event handler *A* will all be executed after *A*, but are
31 not guaranteed to be executed in any particular order amongst themselves.

32 In addition, one event handler can be declared as the *first* handler to be executed in the chain. This
33 handler will *always* be called prior to any other handler, regardless of category, provided the
34 incoming event matches both the specified range and event code. Only one handler can be so
35 designated — attempts to designate additional handlers as *first* will return an error. Deregistration
36 of the declared *first* handler will re-open the position for subsequent assignment.

1 Similarly, one event handler can be declared as the *last* handler to be executed in the chain. This
2 handler will *always* be called after all other handlers have executed, regardless of category,
3 provided the incoming event matches both the specified range and event code. Note that this
4 handler will not be called if the chain is terminated by an earlier handler. Only one handler can be
5 designated as *last* — attempts to designate additional handlers as *last* will return an error.
6 Deregistration of the declared *last* handler will re-open the position for subsequent assignment.

Advice to users

7 Note that the *last* handler is called *after* all registered default handlers that match the specified
8 range of the incoming event unless a handler prior to it terminates the chain. Thus, if the application
9 intends to define a *last* handler, it should ensure that no default handler aborts the process before it.

10 Upon completing its work and prior to returning, each handler *must* call the event handler
11 completion function provided when it was invoked (including a status code plus any information to
12 be passed to later handlers) so that the chain can continue being progressed. PMIx automatically
13 aggregates the status and any results of each handler (as provided in the completion callback) with
14 status from all prior handlers so that each step in the chain has full knowledge of what preceded it.
15 An event handler can terminate all further progress along the chain by passing the
16 `PMIX_EVENT_ACTION_COMPLETE` status to the completion callback function.

17 8.1.1 PMIx_Register_event_handler

18 Summary

19 Register an event handler

20 Format

PMIx v2.0

C

```
21 void  
22 PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,  
23                             pmix_info_t info[], size_t ninfo,  
24                             pmix_notification_fn_t evhdlr,  
25                             pmix_evhdlr_reg_cbfunc_t cbfunc,  
26                             void *cbdata);
```

C

27 **IN** `codes`
28 Array of status codes (array of `pmix_status_t`)
29 **IN** `ncodes`
30 Number of elements in the `codes` array (`size_t`)
31 **IN** `info`
32 Array of info structures (array of handles)

1 **IN** **ninfo**
 2 Number of elements in the *info* array (**size_t**)
 3 **IN** **evhdlr**
 4 Event handler to be called **pmix_notification_fn_t** (function reference)
 5 **IN** **cbfunc**
 6 Callback function **pmix_evhdlr_reg_cbfunc_t** (function reference)
 7 **IN** **cbdata**
 8 Data to be passed to the cbfunc callback function (memory reference)

9 Upon completion, the callback will receive a status based on the following table:

- 10 **PMIX_SUCCESS** The event handler was successfully registered - the event handler identifier is
 11 returned in the callback.
 12 **PMIX_ERR_BAD_PARAM** One or more of the directives provided in the *info* array was
 13 unrecognized.
 14 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support event notification,
 15 or the host SMS does not support notification of the specified event code.

16 The callback function must not be executed prior to returning from the API, and no events
 17 corresponding to this registration may be delivered prior to the completion of the registration
 18 callback function (*cbfunc*).

▼----- Required Attributes -----▼

19 The following attributes are required to be supported by all PMIx libraries:

- 20 **PMIX_EVENT_HDLR_NAME** "pmix.evname" (**char***)
 21 String name identifying this handler.
 22 **PMIX_EVENT_HDLR_FIRST** "pmix.evfirst" (**bool**)
 23 Invoke this event handler before any other handlers.
 24 **PMIX_EVENT_HDLR_LAST** "pmix.evlast" (**bool**)
 25 Invoke this event handler after all other handlers have been called.
 26 **PMIX_EVENT_HDLR_FIRST_IN_CATEGORY** "pmix.evfirstcat" (**bool**)
 27 Invoke this event handler before any other handlers in this category.
 28 **PMIX_EVENT_HDLR_LAST_IN_CATEGORY** "pmix.evlastcat" (**bool**)
 29 Invoke this event handler after all other handlers in this category have been called.
 30 **PMIX_EVENT_HDLR_BEFORE** "pmix.evbefore" (**char***)
 31 Put this event handler immediately before the one specified in the (**char***) value.
 32 **PMIX_EVENT_HDLR_AFTER** "pmix.evafter" (**char***)
 33 Put this event handler immediately after the one specified in the (**char***) value.
 34 **PMIX_EVENT_HDLR_PREPEND** "pmix.evprepend" (**bool**)
 35 Prepend this handler to the precedence list within its category.

1 **PMIX_EVENT_HDLR_APPEND** "pmix.evappend" (bool)
 2 Append this handler to the precedence list within its category.

3 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrangle" (pmix_data_array_t*)
 4 Array of **pmix_proc_t** defining range of event notification.

5 **PMIX_RANGE** "pmix.range" (pmix_data_range_t)
 6 Value for calls to publish/lookup/unpublish or for monitoring event notifications.

7 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (void *)
 8 Object to be returned whenever the registered callback function **cbfunc** is invoked. The
 9 object will only be returned to the process that registered it.

11 Host environments that implement support for PMIx event notification are required to support the
 12 following attributes:

13 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)
 14 The single process that was affected.

15 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (pmix_data_array_t*)
 16 Array of **pmix_proc_t** defining affected processes.



▼----- Optional Attributes -----▼

17 Host environments that support PMIx event notification *may* offer notifications for environmental
 18 events impacting the job and for SMS events relating to the job. The following attributes are
 19 optional for host environments that support this operation:

20 **PMIX_EVENT_TERMINATE_SESSION** "pmix.evterm.sess" (bool)
 21 The RM intends to terminate this session.

22 **PMIX_EVENT_TERMINATE_JOB** "pmix.evterm.job" (bool)
 23 The RM intends to terminate this job.

24 **PMIX_EVENT_TERMINATE_NODE** "pmix.evterm.node" (bool)
 25 The RM intends to terminate all processes on this node.

26 **PMIX_EVENT_TERMINATE_PROC** "pmix.evterm.proc" (bool)
 27 The RM intends to terminate just this process.

28 **PMIX_EVENT_ACTION_TIMEOUT** "pmix.evtimeout" (int)
 29 The time in seconds before the RM will execute error response.

30 **PMIX_EVENT_SILENT_TERMINATION** "pmix.evsilentterm" (bool)
 31 Do not generate an event when this job normally terminates.



Description

Register an event handler to report events. Note that the codes being registered do *not* need to be PMIx error constants — any integer value can be registered. This allows for registration of non-PMIx events such as those defined by a particular SMS vendor or by an application itself.

Advice to users

In order to avoid potential conflicts, users are advised to only define codes that lie outside the range of the PMIx standard's error codes. Thus, SMS vendors and application developers should constrain their definitions to positive values or negative values beyond the `PMIX_EXTERNAL_ERR_BASE` boundary.

Advice to users

As previously stated, upon completing its work, and prior to returning, each handler *must* call the event handler completion function provided when it was invoked (including a status code plus any information to be passed to later handlers) so that the chain can continue being progressed. An event handler can terminate all further progress along the chain by passing the `PMIX_EVENT_ACTION_COMPLETE` status to the completion callback function. Note that the parameters passed to the event handler (e.g., the *info* and *results* arrays) will cease to be valid once the completion function has been called - thus, any information in the incoming parameters that will be referenced following the call to the completion function must be copied.

8.1.2 PMIx_Deregister_event_handler

Summary

Deregister an event handler.

Format

C

```

1 void
2 PMIx_Deregister_event_handler(size_t evhdlr_ref,
3                               pmix_op_cbfunc_t cbfunc,
4                               void *cbdata);
5

```

C

IN evhdlr_refEvent handler ID returned by registration (**size_t**)**IN cbfunc**Callback function to be executed upon completion of operation **pmix_op_cbfunc_t** (function reference)**IN cbdata**

Data to be passed to the cbfunc callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS** , indicating that the request is being processed - result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will *not* be called

If the provided cbfunc is called to confirm removal of the designated handler, the returned status code will be one of the following:

PMIX_SUCCESS The event handler was successfully deregistered.**PMIX_ERR_BAD_PARAM** The provided *evhdlr_ref* was unrecognized.**PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support event notification.**Description**Deregister an event handler. Note that no events corresponding to the referenced registration may be delivered following completion of the deregistration operation (either return from the API with **PMIX_OPERATION_SUCCEEDED** or execution of the *cbfunc*).**8.1.3 PMIx_Notify_event****Summary**

Report an event for notification via any registered event handler.

1
PMIx v2.0

Format

C

```
2 pmix_status_t
3 PMIx_Notify_event (pmix_status_t status,
4                   const pmix_proc_t *source,
5                   pmix_data_range_t range,
6                   pmix_info_t info[], size_t ninfo,
7                   pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- 8 **IN status**
9 Status code of the event ([pmix_status_t](#))
- 10 **IN source**
11 Pointer to a [pmix_proc_t](#) identifying the original reporter of the event (handle)
- 12 **IN range**
13 Range across which this notification shall be delivered ([pmix_data_range_t](#))
- 14 **IN info**
15 Array of [pmix_info_t](#) structures containing any further info provided by the originator of
16 the event (array of handles)
- 17 **IN ninfo**
18 Number of elements in the *info* array ([size_t](#))
- 19 **IN cbfunc**
20 Callback function to be executed upon completion of operation [pmix_op_cbfunc_t](#)
21 (function reference)
- 22 **IN cbdata**
23 Data to be passed to the cbfunc callback function (memory reference)

24 Returns one of the following:

- 25 **PMIX_SUCCESS** The notification request is valid and is being processed. The callback function
26 will be called when the process-local operation is complete and will provide the resulting
27 status of that operation. Note that this does *not* reflect the success or failure of delivering the
28 event to any recipients. The callback function must not be executed prior to returning from the
29 API.
- 30 **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
31 returned *success* - the *cbfunc* will *not* be called
- 32 **PMIX_ERR_BAD_PARAM** The request contains at least one incorrect entry that prevents it from
33 being processed. The callback function will *not* be called.
- 34 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support event notification,
35 or in the case of a PMIx server calling the API, the range extended beyond the local node and
36 the host SMS environment does not support event notification. The callback function will *not*
37 be called.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_EVENT_NON_DEFAULT "pmix.evnondef" (bool)

Event is not to be delivered to default event handlers.

PMIX_EVENT_CUSTOM_RANGE "pmix.evrage" (pmix_data_array_t*)

Array of `pmix_proc_t` defining range of event notification.

Host environments that implement support for PMIx event notification are required to provide the following attributes for all events generated by the environment:

PMIX_EVENT_AFFECTED_PROC "pmix.evproc" (pmix_proc_t)

The single process that was affected.

PMIX_EVENT_AFFECTED_PROCS "pmix.evaffected" (pmix_data_array_t*)

Array of `pmix_proc_t` defining affected processes.

Description

Report an event for notification via any registered event handler. This function can be called by any PMIx process, including application processes, PMIx servers, and SMS elements. The PMIx server calls this API to report events it detected itself so that the host SMS daemon distribute and handle them, and to pass events given to it by its host down to any attached client processes for processing. Examples might include notification of the failure of another process, detection of an impending node failure due to rising temperatures, or an intent to preempt the application. Events may be locally generated or come from anywhere in the system.

Host SMS daemons call the API to pass events down to its embedded PMIx server both for transmittal to local client processes and for the server's own internal processing.

Client application processes can call this function to notify the SMS and/or other application processes of an event it encountered. Note that processes are not constrained to report status values defined in the official PMIx standard — any integer value can be used. Thus, applications are free to define their own internal events and use the notification system for their own internal purposes.

Advice to users

The callback function will be called upon completion of the `notify_event` function's actions. At that time, any messages required for executing the operation (e.g., to send the notification to the local PMIx server) will have been queued, but may not yet have been transmitted. The caller is required to maintain the input data until the callback function has been executed — the sole purpose of the callback function is to indicate when the input data is no longer required.

CHAPTER 9

Data Packing and Unpacking

1 PMIx intentionally does not include support for internode communications in the standard, instead
2 relying on its host SMS environment to transfer any needed data and/or requests between nodes.
3 These operations frequently involve PMIx-defined public data structures that include binary data.
4 Many HPC clusters are homogeneous, and so transferring the structures can be done rather simply.
5 However, greater effort is required in heterogeneous environments to ensure binary data is correctly
6 transferred. PMIx buffer manipulation functions are provided for this purpose via standardized
7 interfaces to ease adoption.

8 9.1 Data Buffer Type

9 The `pmix_data_buffer_t` structure describes a data buffer used for packing and unpacking.

PMIx v2.0

```
10 typedef struct pmix_data_buffer {  
11     /** Start of my memory */  
12     char *base_ptr;  
13     /** Where the next data will be packed to  
14         (within the allocated memory starting  
15         at base_ptr) */  
16     char *pack_ptr;  
17     /** Where the next data will be unpacked  
18         from (within the allocated memory  
19         starting as base_ptr) */  
20     char *unpack_ptr;  
21     /** Number of bytes allocated (starting  
22         at base_ptr) */  
23     size_t bytes_allocated;  
24     /** Number of bytes used by the buffer  
25         (i.e., amount of data -- including  
26         overhead -- packed in the buffer) */  
27     size_t bytes_used;  
28 } pmix_data_buffer_t;
```

1 9.2 Support Macros

2 PMIx provides a set of convenience macros for creating, initiating, and releasing data buffers.

3 9.2.1 PMIX_DATA_BUFFER_CREATE

4 Summary

5 Allocate memory for a `pmix_data_buffer_t` object and initialize it

6 Format

PMIx v2.0

▼  _____ ▼

7 `PMIX_DATA_BUFFER_CREATE(buffer);`

▲  _____ ▲

8 OUT `buffer`

9 Variable to be assigned the pointer to the allocated `pmix_data_buffer_t` (handle)

10 Description

11 This macro uses *calloc* to allocate memory for the buffer and initialize all fields in it

12 9.2.2 PMIX_DATA_BUFFER_RELEASE

13 Summary

14 Free a `pmix_data_buffer_t` object and the data it contains

15 Format

PMIx v2.0

▼  _____ ▼

16 `PMIX_DATA_BUFFER_RELEASE(buffer);`

▲  _____ ▲

17 IN `buffer`

18 Pointer to the `pmix_data_buffer_t` to be released (handle)

19 Description

20 Free's the data contained in the buffer, and then free's the buffer itself

21 9.2.3 PMIX_DATA_BUFFER_CONSTRUCT

22 Summary

23 Initialize a statically declared `pmix_data_buffer_t` object

1 **Format**
PMIx v2.0 ▼ C

2 **PMIX_DATA_BUFFER_CONSTRUCT**(buffer);
▲ C

3 **IN** **buffer**
4 Pointer to the allocated **pmix_data_buffer_t** that is to be initialized (handle)

5 **Description**
6 Initialize a pre-allocated buffer object

7 9.2.4 PMIX_DATA_BUFFER_DESTRUCT

8 **Summary**
9 Release the data contained in a **pmix_data_buffer_t** object

10 **Format**
PMIx v2.0 ▼ C

11 **PMIX_DATA_BUFFER_DESTRUCT**(buffer);
▲ C

12 **IN** **buffer**
13 Pointer to the **pmix_data_buffer_t** whose data is to be released (handle)

14 **Description**
15 Free's the data contained in a **pmix_data_buffer_t** object

16 9.2.5 PMIX_DATA_BUFFER_LOAD

17 **Summary**
18 Load a blob into a **pmix_data_buffer_t** object

19 **Format**
PMIx v2.0 ▼ C

20 **PMIX_DATA_BUFFER_LOAD**(buffer, data, size);
▲ C

21 **IN** **buffer**
22 Pointer to a pre-allocated **pmix_data_buffer_t** (handle)

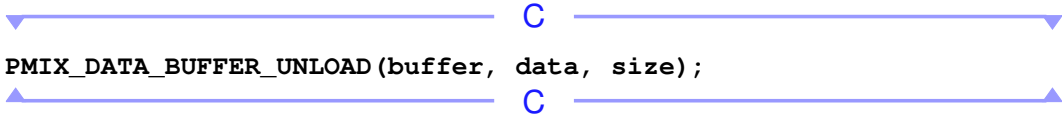
23 **IN** **data**
24 Pointer to a blob (**char***)

25 **IN** **size**
26 Number of bytes in the blob **size_t**

1 **Description**
2 Load the given data into the provided `pmix_data_buffer_t` object, usually done in
3 preparation for unpacking the provided data. Note that the data is *not* copied into the buffer - thus,
4 the blob must not be released until after operations on the buffer have completed.

5 **9.2.6 PMIX_DATA_BUFFER_UNLOAD**

6 **Summary**
7 Unload the data from a `pmix_data_buffer_t` object

8 **Format**
9 
 PMIx v2.0 `PMIX_DATA_BUFFER_UNLOAD(buffer, data, size);`

10 **IN** `buffer`
11 Pointer to the `pmix_data_buffer_t` whose data is to be extracted (handle)
12 **OUT** `data`
13 Variable to be assigned the pointer to the extracted blob (`void*`)
14 **OUT** `size`
15 Variable to be assigned the number of bytes in the blob `size_t`

16 **Description**
17 Extract the data in a buffer, assigning the pointer to the data (and the number of bytes in the blob) to
18 the provided variables, usually done to transmit the blob to a remote process for unpacking. The
19 buffer's internal pointer will be set to NULL to protect the data upon buffer destruct or release -
20 thus, the user is responsible for releasing the blob when done with it.

21 **9.3 General Routines**

22 The following routines are provided to support internode transfers in heterogeneous environments.

23 **9.3.1 PMIx_Data_pack**

24 **Summary**
25 Pack one or more values of a specified type into a buffer, usually for transmission to another process

1
PMIx v2.0

Format

C

```

2 pmix_status_t
3 PMIx_Data_pack(const pmix_proc_t *target,
4               pmix_data_buffer_t *buffer,
5               void *src, int32_t num_vals,
6               pmix_data_type_t type);

```

C

- 7 **IN target**
8 Pointer to a [pmix_proc_t](#) containing the nspace/rank of the process that will be unpacking
9 the final buffer. A NULL value may be used to indicate that the target is based on the same
10 PMIx version as the caller. Note that only the target's nspace is relevant. (handle)
- 11 **IN buffer**
12 Pointer to a [pmix_data_buffer_t](#) where the packed data is to be stored (handle)
- 13 **IN src**
14 Pointer to a location where the data resides. Strings are to be passed as (char **) — i.e., the
15 caller must pass the address of the pointer to the string as the (void*). This allows the caller to
16 pass multiple strings in a single call. (memory reference)
- 17 **IN num_vals**
18 Number of elements pointed to by the *src* pointer. A string value is counted as a single value
19 regardless of length. The values must be contiguous in memory. Arrays of pointers (e.g.,
20 string arrays) should be contiguous, although the data pointed to need not be contiguous
21 across array entries.([int32_t](#))
- 22 **IN type**
23 The type of the data to be packed ([pmix_data_type_t](#))

24 Returns one of the following:

- 25 [PMIX_SUCCESS](#) The data has been packed as requested
- 26 [PMIX_ERR_NOT_SUPPORTED](#) The PMIx implementation does not support this function.
- 27 [PMIX_ERR_BAD_PARAM](#) The provided buffer or src is **NULL**
- 28 [PMIX_ERR_UNKNOWN_DATA_TYPE](#) The specified data type is not known to this
29 implementation
- 30 [PMIX_ERR_OUT_OF_RESOURCE](#) Not enough memory to support the operation
- 31 [PMIX_ERROR](#) General error

Description

33 The pack function packs one or more values of a specified type into the specified buffer. The buffer
34 must have already been initialized via the [PMIX_DATA_BUFFER_CREATE](#) or
35 [PMIX_DATA_BUFFER_CONSTRUCT](#) macros — otherwise, [PMIx_Data_pack](#) will return an
36 error. Providing an unsupported type flag will likewise be reported as an error.

37 Note that any data to be packed that is not hard type cast (i.e., not type cast to a specific size) may
38 lose precision when unpacked by a non-homogeneous recipient. The [PMIx_Data_pack](#) function

1 will do its best to deal with heterogeneity issues between the packer and unpacker in such cases.
2 Sending a number larger than can be handled by the recipient will return an error code (generated
3 upon unpacking) — the error cannot be detected during packing.

4 The namespace of the intended recipient of the packed buffer (i.e., the process that will be
5 unpacking it) is used solely to resolve any data type differences between PMIx versions. The
6 recipient must, therefore, be known to the user prior to calling the pack function so that the PMIx
7 library is aware of the version the recipient is using. Note that all processes in a given namespace
8 are *required* to use the same PMIx version — thus, the caller must only know at least one process
9 from the target’s namespace.

10 9.3.2 PMIx_Data_unpack

11 Summary

12 Unpack values from a `pmix_data_buffer_t`

13 Format

PMIx v2.0

C

```
14 pmix_status_t  
15 PMIx_Data_unpack(const pmix_proc_t *source,  
16                 pmix_data_buffer_t *buffer, void *dest,  
17                 int32_t *max_num_values,  
18                 pmix_data_type_t type);  
19
```

C

20 IN source

21 Pointer to a `pmix_proc_t` structure containing the nspace/rank of the process that packed
22 the provided buffer. A NULL value may be used to indicate that the source is based on the
23 same PMIx version as the caller. Note that only the source’s nspace is relevant. (handle)

24 IN buffer

25 A pointer to the buffer from which the value will be extracted. (handle)

26 INOUT dest

27 A pointer to the memory location into which the data is to be stored. Note that these values
28 will be stored contiguously in memory. For strings, this pointer must be to (`char**`) to provide
29 a means of supporting multiple string operations. The unpack function will allocate memory
30 for each string in the array - the caller must only provide adequate memory for the array of
31 pointers. (`void*`)

32 INOUT max_num_values

33 The number of values to be unpacked — upon completion, the parameter will be set to the
34 actual number of values unpacked. In most cases, this should match the maximum number
35 provided in the parameters — but in no case will it exceed the value of this parameter. Note
36 that unpacking fewer values than are actually available will leave the buffer in an ununpackable
37 state — the function will return an error code to warn of this condition. (`int32_t`)

IN type

The type of the data to be unpacked — must be one of the PMI_x defined data types (`pmix_data_type_t`)

Returns one of the following:

`PMIX_SUCCESS` The data has been unpacked as requested

`PMIX_ERR_NOT_SUPPORTED` The PMI_x implementation does not support this function.

`PMIX_ERR_BAD_PARAM` The provided buffer or dest is **NULL**

`PMIX_ERR_UNKNOWN_DATA_TYPE` The specified data type is not known to this implementation

`PMIX_ERR_OUT_OF_RESOURCE` Not enough memory to support the operation

`PMIX_ERROR` General error

Description

The unpack function unpacks the next value (or values) of a specified type from the given buffer. The buffer must have already been initialized via an `PMIX_DATA_BUFFER_CREATE` or `PMIX_DATA_BUFFER_CONSTRUCT` call (and assumedly filled with some data) — otherwise, the `unpack_value` function will return an error. Providing an unsupported type flag will likewise be reported as an error, as will specifying a data type that *does not* match the type of the next item in the buffer. An attempt to read beyond the end of the stored data held in the buffer will also return an error.

NOTE: it is possible for the buffer to be corrupted and that PMI_x will *think* there is a proper variable type at the beginning of an unpack region — but that the value is bogus (e.g., just a byte field in a string array that so happens to have a value that matches the specified data type flag). Therefore, the data type error check is *not* completely safe.

Unpacking values is a "nondestructive" process — i.e., the values are not removed from the buffer. It is therefore possible for the caller to re-unpack a value from the same buffer by resetting the `unpack_ptr`.

Warning: The caller is responsible for providing adequate memory storage for the requested data. The user must provide a parameter indicating the maximum number of values that can be unpacked into the allocated memory. If more values exist in the buffer than can fit into the memory storage, then the function will unpack what it can fit into that location and return an error code indicating that the buffer was only partially unpacked.

Note that any data that was not hard type cast (i.e., not type cast to a specific size) when packed may lose precision when unpacked by a non-homogeneous recipient. PMI_x will do its best to deal with heterogeneity issues between the packer and unpacker in such cases. Sending a number larger than can be handled by the recipient will return an error code generated upon unpacking — these errors cannot be detected during packing.

The namespace of the process that packed the buffer is used solely to resolve any data type differences between PMI_x versions. The packer must, therefore, be known to the user prior to calling the `pack` function so that the PMI_x library is aware of the version the packer is using. Note

1 that all processes in a given namespace are *required* to use the same PMIx version — thus, the
2 caller must only know at least one process from the packer’s namespace.

3 9.3.3 PMIx_Data_copy

4 Summary

5 Copy a data value from one location to another.

6 Format

PMIx v2.0

C

```
7 pmix_status_t  
8 PMIx_Data_copy(void **dest, void *src,  
9 pmix_data_type_t type);
```

C

10 **IN dest**

11 The address of a pointer into which the address of the resulting data is to be stored. (**void****)

12 **IN src**

13 A pointer to the memory location from which the data is to be copied (handle)

14 **IN type**

15 The type of the data to be copied — must be one of the PMIx defined data types. (

16 [pmix_data_type_t](#))

17 Returns one of the following:

18 [PMIX_SUCCESS](#) The data has been copied as requested

19 [PMIX_ERR_NOT_SUPPORTED](#) The PMIx implementation does not support this function.

20 [PMIX_ERR_BAD_PARAM](#) The provided src or dest is **NULL**

21 [PMIX_ERR_UNKNOWN_DATA_TYPE](#) The specified data type is not known to this
22 implementation

23 [PMIX_ERR_OUT_OF_RESOURCE](#) Not enough memory to support the operation

24 [PMIX_ERROR](#) General error

25 Description

26 Since registered data types can be complex structures, the system needs some way to know how to
27 copy the data from one location to another (e.g., for storage in the registry). This function, which
28 can call other copy functions to build up complex data types, defines the method for making a copy
29 of the specified data type.

30 9.3.4 PMIx_Data_print

31 Summary

32 Pretty-print a data value.

1
PMIx v2.0

Format

C

```
2 pmix_status_t
3 PMIx_Data_print(char **output, char *prefix,
4                 void *src, pmix_data_type_t type);
```

C

- 5 **IN output**
6 The address of a pointer into which the address of the resulting output is to be stored.
7 (**char****)
- 8 **IN prefix**
9 String to be prepended to the resulting output (**char***)
- 10 **IN src**
11 A pointer to the memory location of the data value to be printed (handle)
- 12 **IN type**
13 The type of the data value to be printed — must be one of the PMIx defined data types. (
14 [pmix_data_type_t](#))

15 Returns one of the following:

- 16 [PMIX_SUCCESS](#) The data has been printed as requested
- 17 [PMIX_ERR_BAD_PARAM](#) The provided data type is not recognized.
- 18 [PMIX_ERR_NOT_SUPPORTED](#) The PMIx implementation does not support this function.

Description

20 Since registered data types can be complex structures, the system needs some way to know how to
21 print them (i.e., convert them to a string representation). Primarily for debug purposes.

22 9.3.5 PMIx_Data_copy_payload

23 Summary

24 Copy a payload from one buffer to another

25 Format

PMIx v2.0

C

```
1 pmix_status_t  
2 PMIx_Data_copy_payload(pmix_data_buffer_t *dest,  
3 pmix_data_buffer_t *src);
```

C

```
4 IN dest  
5     Pointer to the destination pmix_data_buffer_t (handle)  
6 IN src  
7     Pointer to the source pmix_data_buffer_t (handle)
```

8 Returns one of the following:

```
9 PMIX_SUCCESS The data has been copied as requested  
10 PMIX_ERR_BAD_PARAM The src and dest pmix_data_buffer_t types do not match  
11 PMIX_ERR_NOT_SUPPORTED The PMIx implementation does not support this function.
```

12 Description

13 This function will append a copy of the payload in one buffer into another buffer. Note that this is
14 *not* a destructive procedure — the source buffer’s payload will remain intact, as will any pre-existing
15 payload in the destination’s buffer. Only the unpacked portion of the source payload will be copied.

CHAPTER 10

Security

1 PMIx utilizes a multi-layered approach toward security that differs for client versus tool processes.
2 *Client* processes (i.e., processes started by the host environment) must be preregistered with the
3 PMIx server library via the `PMIx_server_register_client` API before they are spawned.
4 This API requires that you pass the expected uid/gid of the client process.

5 When the client attempts to connect to the PMIx server, the server uses available standard
6 Operating System (OS) methods to determine the effective uid/gid of the process requesting the
7 connection. PMIx implementations shall not rely on any values reported by the client process itself
8 as that would be unsafe. The effective uid/gid reported by the OS is compared to the values
9 provided by the host during registration - if they don't match, the PMIx server is required to drop
10 the connection request. This ensures that the PMIx server does not allow connection from a client
11 that doesn't at least meet some minimal security requirement.

12 Once the requesting client passes the initial test, the PMIx server can, at the choice of the
13 implementor, perform additional security checks. This may involve a variety of methods such as
14 exchange of a system-provided key or credential. At the conclusion of that process, the PMIx server
15 reports the client connection request to the host via the
16 `pmix_server_client_connected_fn_t` interface. The host may then perform any
17 additional checks and operations before responding with either `PMIX_SUCCESS` to indicate that
18 the connection is approved, or a PMIx error constant indicating that the connection request is
19 refused. In this latter case, the PMIx server is required to drop the connection.

20 Tools started by the host environment are classed as a subgroup of client processes and follow the
21 client process procedure. However, tools that are not started by the host environment must be
22 handled differently as registration information is not available prior to the connection request. In
23 these cases, the PMIx server library is required to use available standard OS methods to get the
24 effective uid/gid and report them upwards as part of invoking the
25 `pmix_server_tool_connection_fn_t` interface, deferring initial security screening to
26 the host. It is recognized that this may represent a security risk - for this reason, PMIx server
27 libraries must not enable tool connections by default. Instead, the host has to explicitly enable them
28 via the `PMIX_SERVER_TOOL_SUPPORT` attribute, thus recognizing the associated risk. Once
29 the host has completed its authentication procedure, it again informs the PMIx server of the result.

30 Applications and tools often interact with the host environment in ways that require security beyond
31 just verifying the user's identity - e.g., access to that user's relevant authorizations. This is
32 particularly important when tools connect directly to a system-level PMIx server that may be
33 operating at a privileged level. A variety of system management software packages provide
34 authorization services, but the lack of standardized interfaces makes portability problematic.

1 This section defines two PMIx client-side APIs for this purpose. These are most likely to be used
2 by user-space applications/tools, but are not restricted to that realm.

3 10.1 Obtaining Credentials

4 The API for obtaining a credential is a non-blocking operation since the host environment may have
5 to contact a remote credential service. The definition takes into account the potential that the
6 returned credential could be sent via some mechanism to another application that resides in an
7 environment using a different security mechanism. Thus, provision is made for the system to return
8 additional information (e.g., the identity of the issuing agent) outside of the credential itself and
9 visible to the application.

10 10.1.1 PMIx_Get_credential

11 Summary

12 Request a credential from the PMIx server library or the host environment

13 Format

PMIx v3.0

```
14 pmix_status_t  
15 PMIx_Get_credential(const pmix_info_t info[], size_t ninfo,  
16 pmix_credential_cbfunc_t cbfunc, void *cbdata)  
17 C
```

- 17 **IN info**
18 Array of **pmix_info_t** structures (array of handles)
- 19 **IN ninfo**
20 Number of elements in the *info* array (**size_t**)
- 21 **IN cbfunc**
22 Callback function to return credential (**pmix_credential_cbfunc_t** function
23 reference)
- 24 **IN cbdata**
25 Data to be passed to the callback function (memory reference)

26 Returns one of the following:

- 27 • **PMIX_SUCCESS**, indicating that the request has been communicated to the local PMIx server -
28 result will be returned in the provided *cbfunc*
- 29 • a PMIx error constant indicating either an error in the input or that the request is unsupported -
30 the *cbfunc* will *not* be called

Required Attributes

1 PMIx libraries that choose not to support this operation *must* return
2 **PMIX_ERR_NOT_SUPPORTED** when the function is called.

3 There are no required attributes for this API. Note that implementations may choose to internally
4 execute integration for some security environments (e.g., directly contacting a *munge* server).

5 Implementations that support the operation but cannot directly process the client's request must
6 pass any attributes that are provided by the client to the host environment for processing. In
7 addition, the following attributes are required to be included in the *info* array passed from the PMIx
8 library to the host environment:

9 **PMIX_USERID** "pmix.euid" (uint32_t)
10 Effective user id.

11 **PMIX_GRPID** "pmix.egid" (uint32_t)
12 Effective group id.

Optional Attributes

13 The following attributes are optional for host environments that support this operation:

14 **PMIX_TIMEOUT** "pmix.timeout" (int)
15 Time in seconds before the specified operation should time out (0 indicating infinite) in
16 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
17 the target process from ever exposing its data.

Advice to PMIx library implementers

18 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
19 environment due to race condition considerations between completion of the operation versus
20 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
21 directly in the PMIx server library must take care to resolve the race condition and should avoid
22 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
23 created.

Description

24 Request a credential from the PMIx server library or the host environment
25

1 10.2 Validating Credentials

2 The API for validating a credential is a non-blocking operation since the host environment may
3 have to contact a remote credential service. Provision is made for the system to return additional
4 information regarding possible authorization limitations beyond simple authentication.

5 10.2.1 PMIx_Validate_credential

6 Summary

7 Request validation of a credential by the PMIx server library or the host environment

8 Format

PMIx v3.0

C

```
9 pmix_status_t  
10 PMIx_Validate_credential(const pmix_byte_object_t *cred,  
11                          const pmix_info_t info[], size_t ninfo,  
12                          pmix_validation_cbfunc_t cbfunc,  
13                          void *cbdata)
```

C

14 IN cred

15 Pointer to [pmix_byte_object_t](#) containing the credential (handle)

16 IN info

17 Array of [pmix_info_t](#) structures (array of handles)

18 IN ninfo

19 Number of elements in the *info* array ([size_t](#))

20 IN cbfunc

21 Callback function to return result ([pmix_validation_cbfunc_t](#) function reference)

22 IN cbdata

23 Data to be passed to the callback function (memory reference)

24 Returns one of the following:

- 25 • [PMIX_SUCCESS](#) , indicating that the request has been communicated to the local PMIx server -
26 result will be returned in the provided *cbfunc*
- 27 • a PMIx error constant indicating either an error in the input or that the request is unsupported -
28 the *cbfunc* will *not* be called

Required Attributes

1 PMIx libraries that choose not to support this operation *must* return
2 **PMIX_ERR_NOT_SUPPORTED** when the function is called.

3 There are no required attributes for this API. Note that implementations may choose to internally
4 execute integration for some security environments (e.g., directly contacting a *munge* server).

5 Implementations that support the operation but cannot directly process the client's request must
6 pass any attributes that are provided by the client to the host environment for processing. In
7 addition, the following attributes are required to be included in the *info* array passed from the PMIx
8 library to the host environment:

9 **PMIX_USERID** "pmix.euid" (uint32_t)
10 Effective user id.

11 **PMIX_GRPID** "pmix.egid" (uint32_t)
12 Effective group id.

Optional Attributes

13 The following attributes are optional for host environments that support this operation:

14 **PMIX_TIMEOUT** "pmix.timeout" (int)
15 Time in seconds before the specified operation should time out (0 indicating infinite) in
16 error. The timeout parameter can help avoid "hangs" due to programming errors that prevent
17 the target process from ever exposing its data.

Advice to PMIx library implementers

18 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
19 environment due to race condition considerations between completion of the operation versus
20 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
21 directly in the PMIx server library must take care to resolve the race condition and should avoid
22 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
23 created.

Description

24 Request validation of a credential by the PMIx server library or the host environment.
25

CHAPTER 11

Server-Specific Interfaces

1 The RM daemon that hosts the PMIx server library interacts with that library in two distinct
2 manners. First, PMIx provides a set of APIs by which the host can request specific services from its
3 library. This includes generating regular expressions, registering information to be passed to client
4 processes, and requesting information on behalf of a remote process. Note that the host always has
5 access to all PMIx client APIs - the functions listed below are in addition to those available to a
6 PMIx client.

7 Second, the host can provide a set of callback functions by which the PMIx server library can pass
8 requests upward for servicing by the host. These include notifications of client connection and
9 finalize, as well as requests by clients for information and/or services that the PMIx server library
10 does not itself provide.

11 11.1 Server Support Functions

12 The following APIs allow the RM daemon that hosts the PMIx server library to request specific
13 services from the PMIx library.

14 11.1.1 PMIx_generate_regex

15 Summary

16 Generate a compressed representation of the input string.

17 Format

PMIx v1.0

18 `pmix_status_t`

19 `PMIx_generate_regex(const char *input, char **output)`

20 **IN** `input`

String to process (string)

22 **OUT** `output`

23 Compressed representation of *input* (array of bytes)

24 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Given a comma-separated list of *input* values, generate a reduced size representation of the input that can be passed down to the PMIx server library's `PMIx_server_register_nspace` API for parsing. The order of the individual values in the *input* string is preserved across the operation. The caller is responsible for releasing the returned data.

The precise compressed representations will be implementation specific. However, all PMIx implementations are required to include a **NULL**-terminated string in the output representation that can be printed for diagnostic purposes.

Advice to PMIx server hosts

The returned representation may be an arbitrary array of bytes as opposed to a valid **NULL**-terminated string. However, the method used to generate the representation shall be identified with a colon-delimited string at the beginning of the output. For example, an output starting with "`pmix:\0`" might indicate that the representation is a PMIx-defined regular expression represented as a **NULL**-terminated string following the "`pmix:\0`" prefix. In contrast, an output starting with "`blob:\0`" might indicate a compressed binary array follows the prefix.

Communicating the resulting output should be done by first packing the returned expression using the `PMIx_Data_pack`, declaring the input to be of type `PMIX_REGEX`, and then obtaining the resulting blob to be communicated using the `PMIX_DATA_BUFFER_UNLOAD` macro. The reciprocal method can be used on the remote end prior to passing the regex into `PMIx_server_register_nspace`. The pack/unpack routines will ensure proper handling of the data based on the regex prefix.

11.1.2 PMIx_generate_ppn

Summary

Generate a compressed representation of the input identifying the processes on each node.

Format

PMIx v1.0

```
pmix_status_t PMIx_generate_ppn(const char *input, char **ppn)
```

IN `input`

String to process (string)

OUT `ppn`

Compressed representation of *input* (array of bytes)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

The input shall consist of a semicolon-separated list of ranges representing the ranks of processes on each node of the job - e.g., "1-4;2-5;8,10,11,12;6,7,9". Each field of the input must correspond to the node name provided at that position in the input to `PMIx_generate_regex`. Thus, in the example, ranks 1-4 would be located on the first node of the comma-separated list of names provided to `PMIx_generate_regex`, and ranks 2-5 would be on the second name in the list.

Advice to PMIx server hosts

The returned representation may be an arbitrary array of bytes as opposed to a valid NULL-terminated string. However, the method used to generate the representation shall be identified with a colon-delimited string at the beginning of the output. For example, an output starting with "`pmix:`" indicates that the representation is a PMIx-defined regular expression represented as a NULL-terminated string. In contrast, an output starting with "`blob:\0size=1234:`" is a compressed binary array.

Communicating the resulting output should be done by first packing the returned expression using the `PMIx_Data_pack`, declaring the input to be of type `PMIX_REGEX`, and then obtaining the blob to be communicated using the `PMIX_DATA_BUFFER_UNLOAD` macro. The pack/unpack routines will ensure proper handling of the data based on the regex prefix.

11.1.3 PMIx_server_register_namespace

Summary

Setup the data about a particular namespace.

Format

PMIx v1.0

```
pmix_status_t
PMIx_server_register_namespace(const pmix_namespace_t namespace,
                               int nlocalprocs,
                               pmix_info_t info[], size_t ninfo,
                               pmix_op_cbfunc_t cbfunc, void *cbdata)
```

IN namespace

Character array of maximum size `PMIX_MAX_NSLEN` containing the namespace identifier (string)

IN nlocalprocs

number of local processes (integer)

IN info

Array of info structures (array of handles)

1 **IN ninfo**
2 Number of elements in the *info* array (integer)
3 **IN cbfunc**
4 Callback function `pmix_op_cbfunc_t` (function reference)
5 **IN cbdata**
6 Data to be passed to the callback function (memory reference)


7 Returns one of the following:

- 8 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
9 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
10 function prior to returning from the API.
- 11 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
12 returned *success* - the *cbfunc* will not be called
- 13 • a PMIx error constant indicating either an error in the input or that the request was immediately
14 processed and failed - the *cbfunc* will not be called

▼----- Required Attributes -----▼

15 The following attributes are required to be supported by all PMIx libraries:

16 **PMIX_REGISTER_NODATA** "pmix.reg.nodata" (bool)
17 Registration is for this namespace only, do not copy job data - this attribute is not accessed
18 using the `PMIx_Get`



20 Host environments are required to provide the following attributes:

- 21 • for the session containing the given namespace:
 - 22 – **PMIX_UNIV_SIZE** "pmix.univ.size" (uint32_t)
23 Number of allocated slots in a session - each slot may or may not be occupied by an
24 executing process. Note that this attribute is the equivalent to the combination of
25 **PMIX_SESSION_INFO_ARRAY** with the **PMIX_MAX_PROCS** entry in the array - it
26 is included in the Standard for historical reasons.
- 27 • for the given namespace:
 - 28 – **PMIX_JOBID** "pmix.jobid" (char*)
29 Job identifier assigned by the scheduler.
 - 30 – **PMIX_JOB_SIZE** "pmix.job.size" (uint32_t)
31 Total number of processes in this job across all contained applications. Note that this
32 value can be different from **PMIX_MAX_PROCS**. For example, users may choose to
33 subdivide an allocation (running several jobs in parallel within it), and dynamic
34 programming models may support adding and removing processes from a running *job*

1 on-they-fly. In the latter case, PMIx events must be used to notify processes within the
2 job that the job size has changed.

3 - **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)
4 Maximum number of processes that can be executed in this context (session,
5 namespace, application, or node). Typically, this is a constraint imposed by a scheduler
6 or by user settings in a hostfile or other resource description.

7 - **PMIX_NODE_MAP** "pmix.nmap" (char*)
8 Regular expression of nodes - see 11.1.3.1 for an explanation of its generation.

9 - **PMIX_PROC_MAP** "pmix.pmap" (char*)
10 Regular expression describing processes on each node - see 11.1.3.1 for an explanation
11 of its generation.

- for its own node:

12 - **PMIX_LOCAL_SIZE** "pmix.local.size" (uint32_t)
13 Number of processes in this job or application on this node.

14 - **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)
15 Comma-delimited list of ranks on this node within the specified namespace - referenced
16 using **PMIX_RANK_WILDCARD**.

17 - **PMIX_LOCAL_CPUSSETS** "pmix.lcpus" (char*)
18 Colon-delimited cpusets of local peers within the specified namespace - referenced
19 using **PMIX_RANK_WILDCARD**.

- for each process in the given namespace:

20 - **PMIX_RANK** "pmix.rank" (pmix_rank_t)
21 Process rank within the job.

22 - **PMIX_LOCAL_RANK** "pmix.lrank" (uint16_t)
23 Local rank on this node within this job.

24 - **PMIX_NODE_RANK** "pmix.nrank" (uint16_t)
25 Process rank on this node spanning all jobs.


26 - **PMIX_NODEID** "pmix.nodeid" (uint32_t)
27 Node identifier where the specified process is located, expressed as the node's index
28 (beginning at zero) in the array resulting from expansion of the **PMIX_NODE_MAP**
29 regular expression for the **job**

30 If more than one application is included in the namespace, then the host environment is also
31 required to provide the following attributes:

- for each application:

32 - **PMIX_APPNUM** "pmix.appnum" (uint32_t)
33 Application number within the job.

- 1 - **PMIX_APPLDR** "pmix.aldr" (pmix_rank_t)
- 2 Lowest rank in this application within this job - referenced using
- 3 **PMIX_RANK_WILDCARD** .
- 4 - **PMIX_APP_SIZE** "pmix.app.size" (uint32_t)
- 5 Number of processes in this application.
- 6 • for each process:
- 7 - **PMIX_APP_RANK** "pmix.apprank" (pmix_rank_t)
- 8 Process rank within this application.
- 9 - **PMIX_APPNUM** "pmix.appnum" (uint32_t)
- 10 Application number within the job.



11 The following attributes may be provided by host environments:

- 12 • for the session containing the given namespace:
- 13 - **PMIX_SESSION_ID** "pmix.session.id" (uint32_t)
- 14 Session identifier - referenced using **PMIX_RANK_WILDCARD** .
- 15 • for the given namespace:
- 16 - **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (char*)
- 17 Name of the namespace to use for this PMIx server.
- 18 - **PMIX_SERVER_RANK** "pmix.srv.rank" (pmix_rank_t)
- 19 Rank of this PMIx server
- 20 - **PMIX_NPROC_OFFSET** "pmix.offset" (pmix_rank_t)
- 21 Starting global rank of this job - referenced using **PMIX_RANK_WILDCARD** .
- 22 - **PMIX_ALLOCATED_NODELIST** "pmix.alist" (char*)
- 23 Comma-delimited list of all nodes in this allocation regardless of whether or not they
- 24 currently host processes - referenced using **PMIX_RANK_WILDCARD** .
- 25 - **PMIX_JOB_NUM_APPS** "pmix.job.napps" (uint32_t)
- 26 Number of applications in this job.
- 27 - **PMIX_MAPBY** "pmix.mapby" (char*)
- 28 Process mapping policy - when accessed using **PMIx_Get** , use the
- 29 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for
- 30 the provided namespace
- 31 - **PMIX_RANKBY** "pmix.rankby" (char*)

1 Process ranking policy - when accessed using `PMIx_Get`, use the
2 `PMIX_RANK_WILDCARD` value for the rank to discover the ranking algorithm used
3 for the provided namespace

4 - `PMIX_BINDTO` "pmix.bindto" (char*)
5 Process binding policy - when accessed using `PMIx_Get`, use the
6 `PMIX_RANK_WILDCARD` value for the rank to discover the binding policy used for
7 the provided namespace

8 - `PMIX_ANL_MAP` "pmix.anlmap" (char*)
9 Process mapping in Argonne National Laboratory's PMI-1/PMI-2 notation.

10 • for its own node:

11 - `PMIX_AVAIL_PHYS_MEMORY` "pmix.pmem" (uint64_t)
12 Total available physical memory on this node.

13 - `PMIX_HWLOC_XML_V1` "pmix.hwlocxml1" (char*)
14 XML representation of local topology using HWLOC's v1.x format.

15 - `PMIX_HWLOC_XML_V2` "pmix.hwlocxml2" (char*)
16 XML representation of local topology using HWLOC's v2.x format.

17 - `PMIX_LOCALLDR` "pmix.lldr" (pmix_rank_t)
18 Lowest rank on this node within this job - referenced using `PMIX_RANK_WILDCARD`.
19

20 - `PMIX_NODE_SIZE` "pmix.node.size" (uint32_t)
21 Number of processes across all jobs on this node.

22 - `PMIX_LOCAL_PROCS` "pmix.lprocs" (pmix_proc_t array)
23 Array of `pmix_proc_t` of all processes on the specified node - referenced using
24 `PMIX_RANK_WILDCARD`.

25 • for each process in the given namespace:

26 - `PMIX_PROCID` "pmix.procid" (pmix_proc_t)
27 Process identifier

28 - `PMIX_GLOBAL_RANK` "pmix.grank" (pmix_rank_t)
29 Process rank spanning across all jobs in this session.

30 - `PMIX_HOSTNAME` "pmix.hname" (char*)
31 Name of the host where the specified process is running.

32 Attributes not directly provided by the host environment may be derived by the PMIx server library
33 from other required information and included in the data made available to the server library's
34 clients.



Description

Pass job-related information to the PMIx server library for distribution to local client processes.

Advice to PMIx server hosts

Host environments are required to execute this operation prior to starting any local application process within the given namespace.

The PMIx server must register all namespaces that will participate in collective operations with local processes. This means that the server must register a namespace even if it will not host any local processes from within that namespace if any local process of another namespace might at some point perform an operation involving one or more processes from the new namespace. This is necessary so that the collective operation can identify the participants and know when it is locally complete.

The caller must also provide the number of local processes that will be launched within this namespace. This is required for the PMIx server library to correctly handle collectives as a collective operation call can occur before all the local processes have been started.

Advice to users

The number of local processes for any given namespace is generally fixed at the time of application launch. Calls to `PMIx_Spawn` result in processes launched in their own namespace, not that of their parent. However, it is possible for processes to *migrate* to another node via a call to `PMIx_Job_control_nb`, thus resulting in a change to the number of local processes on both the initial node and the node to which the process moved. It is therefore critical that applications not migrate processes without first ensuring that PMIx-based collective operations are not in progress, and that no such operations be initiated until process migration has completed.

1 11.1.3.1 Assembling the registration information

2 The following description is not intended to represent the actual layout of information in a given
3 PMIx library. Instead, it describes how information provided in the *info* parameter of the
4 `PMIx_server_register_nspace` shall be organized for proper processing by a PMIx server
5 library. The ordering of the various information elements is arbitrary - they are presented in a
6 top-down hierarchical form solely for clarity in reading.

▼ Advice to PMIx server hosts ▼

7 Creating the *info* array of data requires knowing in advance the number of elements required for the
8 array. This can be difficult to compute and somewhat fragile in practice. One method for resolving
9 the problem is to create a linked list of objects, each containing a single `pmix_info_t` structure.
10 Allocation and manipulation of the list can then be accomplished using existing standard methods.
11 Upon completion, the final *info* array can be allocated based on the number of elements on the list,
12 and then the values in the list object `pmix_info_t` structures transferred to the corresponding
13 array element utilizing the `PMIX_INFO_XFER` macro.

14 A common building block used in several areas is the construction of a regular expression
15 identifying the nodes involved in that area - e.g., the nodes in a `session` or `job`. PMIx provides
16 several tools to facilitate this operation, beginning by constructing an argv-like array of node
17 names. This array is then passed to the `PMIx_generate_regex` function to create a regular
18 expression parseable by the PMIx server library, as shown below:

▼ C ▼

```
19 char **nodes = NULL;  
20 char *nodelist;  
21 char *regex;  
22 size_t n;  
23 pmix_status_t rc;  
24 pmix_info_t info;  
25  
26 /* loop over an array of nodes, adding each  
27 * name to the array */  
28 for (n=0; n < num_nodes; n++)  
29     /* filter the nodes to ignore those not included  
30     * in the target range (session, job, etc.). In  
31     * this example, all nodes are accepted */  
32     PMIX_ARGV_APPEND(&nodes, node[n]->name);  
33  
34  
35 /* join into a comma-delimited string */  
36 nodelist = PMIX_ARGV_JOIN(nodes, ',');  
37
```

```

1      /* release the array */
2      PMIX_ARGV_FREE(nodes);
3
4      /* generate regex */
5      rc = PMIx_generate_regex(nodelist, &regex);
6
7      /* release list */
8      free(nodelist);
9
10     /* pass the regex as the value to the PMIX_NODE_MAP key */
11     PMIX_INFO_LOAD(&info, PMIX_NODE_MAP, regex, PMIX_STRING);
12     /* release the regex */
13     free(regex);
14

```



15 Changing the filter criteria allows the construction of node maps for any level of information.

16 A similar method is used to construct the map of processes on each node from the namespace being
17 registered. This may be done for each information level of interest (e.g., to identify the process map
18 for the entire **job** or for each **application** in the job) by changing the search criteria. An
19 example is shown below for the case of creating the process map for a **job** :



```

20     char **ndppn;
21     char rank[30];
22     char **ppnarray = NULL;
23     char *ppn;
24     char *localranks;
25     char *regex;
26     size_t n, m;
27     pmix_status_t rc;
28     pmix_info_t info;
29
30     /* loop over an array of nodes */
31     for (n=0; n < num_nodes; n++)
32         /* for each node, construct an array of ranks on that node */
33         ndppn = NULL;
34         for (m=0; m < node[n]->num_procs; m++)
35             /* ignore processes that are not part of the target job */
36             if (!PMIX_CHECK_NAMESPACE(targetjob, node[n]->proc[m].namespace))
37                 continue;
38
39             snprintf(rank, 30, "%d", node[n]->proc[m].rank);
40             PMIX_ARGV_APPEND(&ndppn, rank);

```

```

1
2     /* convert the array into a comma-delimited string of ranks */
3     localranks = PMIX_ARGV_JOIN(ndppn, ',');
4     /* release the local array */
5     PMIX_ARGV_FREE(ndppn);
6     /* add this node's contribution to the overall array */
7     PMIX_ARGV_APPEND(&ppnarray, localranks);
8     /* release the local list */
9     free(localranks);
10
11
12     /* join into a semicolon-delimited string */
13     ppn = PMIX_ARGV_JOIN(ppnarray, ';');
14
15     /* release the array */
16     PMIX_ARGV_FREE(ppnarray);
17
18     /* generate ppn regex */
19     rc = PMIx_generate_ppn(ppn, &regex);
20
21     /* release list */
22     free(ppn);
23
24     /* pass the regex as the value to the PMIX_PROC_MAP key */
25     PMIX_INFO_LOAD(&info, PMIX_PROC_MAP, regex, PMIX_STRING);
26     /* release the regex */
27     free(regex);
28

```

C

29 Note that the [PMIX_NODE_MAP](#) and [PMIX_PROC_MAP](#) attributes are linked in that the order of
30 entries in the process map must match the ordering of nodes in the node map - i.e., there is no
31 provision in the PMIx process map regular expression generator/parser pair supporting an
32 out-of-order node or a node that has no corresponding process map entry (e.g., a node with no
33 processes on it). Armed with these tools, the registration *info* array can be constructed as follows:

- 34 • Session-level information includes all session-specific values. In many cases, only two values (
35 [PMIX_SESSION_ID](#) and [PMIX_UNIV_SIZE](#)) are included in the registration array. Since
36 both of these values are session-specific, they can be specified independently - i.e., in their own
37 [pmix_info_t](#) elements of the *info* array. Alternatively, they can be provided as a
38 [pmix_data_array_t](#) array of [pmix_info_t](#) using the [PMIX_SESSION_INFO_ARRAY](#)
39 attribute and identified by including the [PMIX_SESSION_ID](#) attribute in the array - this is
40 required in cases where non-specific attributes (e.g., [PMIX_NUM_NODES](#) or [PMIX_NODE_MAP](#)

1) are passed to describe aspects of the session. Note that the node map can include nodes not
 2 used by the job being registered as no corresponding process map is specified.

3 The *info* array at this point might look like (where the labels identify the corresponding attribute
 4 - e.g., “Session ID” corresponds to the **PMIX_SESSION_ID** attribute):




Figure 11.1.: Session-level information elements

5 • Job-level information includes all job-specific values such as **PMIX_JOB_SIZE** ,
 6 **PMIX_JOB_NUM_APPS** , and **PMIX_JOBID** . Since each invocation of
 7 **PMIx_server_register_namespace** describes a single *job* , job-specific values can be
 8 specified independently - i.e., in their own **pmix_info_t** elements of the *info* array.
 9 Alternatively, they can be provided as a **pmix_data_array_t** array of **pmix_info_t**
 10 identified by the **PMIX_JOB_INFO_ARRAY** attribute - this is required in cases where
 11 non-specific attributes (e.g., **PMIX_NODE_MAP**) are passed to describe aspects of the job. Note
 12 that since the invocation only involves a single namespace, there is no need to include the
 13 **PMIX_NAMESPACE** attribute in the array.

14 Upon conclusion of this step, the *info* array might look like:

15 Note that in this example, **PMIX_NUM_NODES** is not required as that information is contained
 16 in the **PMIX_NODE_MAP** attribute. Similarly, **PMIX_JOB_SIZE** is not technically required as
 17 that information is contained in the **PMIX_PROC_MAP** when combined with the corresponding
 18 node map - however, there is no issue with including the job size as a separate entry.

19 The example also illustrates the hierarchical use of the **PMIX_NODE_INFO_ARRAY** attribute.
 20 In this case, we have chosen to pass several job-related values for each node - since those values
 21 are non-unique across the job, they must be passed in a node-info container. Note that the choice
 22 of what information to pass into the PMIx server library versus what information to derive from
 23 other values at time of request is left to the host environment. PMIx implementors in turn may, if
 24 they choose, pre-parse registration data to create expanded views (thus enabling faster response
 25 to requests at the expense of memory footprint) or to compress views into tighter representations
 26 (thus trading minimized footprint for longer response times).

27 • Application-level information includes all application-specific values such as **PMIX_APP_SIZE**
 28 and **PMIX_APPLDR** . If the *job* contains only a single *application* , then the
 29 application-specific values can be specified independently - i.e., in their own **pmix_info_t**




Figure 11.2.: Job-level information elements

elements of the *info* array - or as a `pmix_data_array_t` array of `pmix_info_t` using the `PMIX_APP_INFO_ARRAY` attribute and identified by including the `PMIX_APPNUM` attribute in the array. Use of the array format is must in cases where non-specific attributes (e.g., `PMIX_NODE_MAP`) are passed to describe aspects of the application.

However, in the case of a job consisting of multiple applications, all application-specific values for each application must be provided using the `PMIX_APP_INFO_ARRAY` format, each identified by its `PMIX_APPNUM` value.

Upon conclusion of this step, the *info* array might look like that shown in 11.3, assuming there are two applications in the job being registered:

- Process-level information includes an entry for each process in the job being registered, each entry marked with the `PMIX_PROC_DATA` attribute. The `rank` of the process must be the first entry in the array - this provides efficiency when storing the data. Upon conclusion of this step, the *info* array might look like the diagram in 11.4:
- For purposes of this example, node-level information only includes values describing the local node - i.e., it does not include information about other nodes in the job or session. In many cases,




Figure 11.3.: Application-level information elements

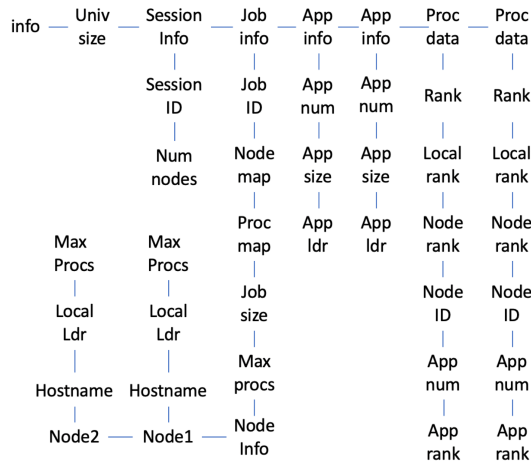


Figure 11.4.: Process-level information elements

- 1 the values included in this level are unique to it and can be specified independently - i.e., in their
- 2 own `pmix_info_t` elements of the `info` array. Alternatively, they can be provided as a
- 3 `pmix_data_array_t` array of `pmix_info_t` using the `PMIX_NODE_INFO_ARRAY`

1 attribute - this is required in cases where non-specific attributes are passed to describe aspects of
2 the node, or where values for multiple nodes are being provided.

3 The node-level information requires two elements that must be constructed in a manner similar to
4 that used for the node map. The **PMIX_LOCAL_PEERS** value is computed based on the
5 processes on the local node, filtered to select those from the job being registered, as shown below
6 using the tools provided by PMIx:

C

```
7 char **ndppn = NULL;
8 char rank[30];
9 char *localranks;
10 size_t m;
11 pmix_info_t info;
12
13 for (m=0; m < mynode->num_procs; m++)
14     /* ignore processes that are not part of the target job */
15     if (!PMIX_CHECK_NAMESPACE(targetjob, mynode->proc[m].nspace))
16         continue;
17
18     snprintf(rank, 30, "%d", mynode->proc[m].rank);
19     PMIX_ARGV_APPEND(&ndppn, rank);
20
21     /* convert the array into a comma-delimited string of ranks */
22     localranks = PMIX_ARGV_JOIN(ndppn, ',');
23     /* release the local array */
24     PMIX_ARGV_FREE(ndppn);
25
26     /* pass the string as the value to the PMIX_LOCAL_PEERS key */
27     PMIX_INFO_LOAD(&info, PMIX_LOCAL_PEERS, localranks, PMIX_STRING);
28     /* release the list */
29     free(localranks);
30
```

C

31 The **PMIX_LOCAL_CPUSSETS** value is constructed in a similar manner. In the provided
32 example, it is assumed that the Hardware Locality (HWLOC) cpuset representation (a
33 comma-delimited string of processor IDs) of the processors assigned to each process has
34 previously been generated and stored on the process description. Thus, the value can be
35 constructed as shown below:

```

1      char **ndcpus = NULL;
2      char *localcpus;
3      size_t m;
4      pmix_info_t info;
5
6      for (m=0; m < mynode->num_procs; m++)
7          /* ignore processes that are not part of the target job */
8          if (!PMIX_CHECK_NAMESPACE(targetjob, mynode->proc[m].nspace))
9              continue;
10
11         PMIX_ARGV_APPEND(&ndcpus, mynode->proc[m].cpuset);
12
13         /* convert the array into a colon-delimited string */
14         localcpus = PMIX_ARGV_JOIN(ndcpus, ':');
15         /* release the local array */
16         PMIX_ARGV_FREE(ndcpus);
17
18         /* pass the string as the value to the PMIX_LOCAL_CPUSSETS key */
19         PMIX_INFO_LOAD(&info, PMIX_LOCAL_CPUSSETS, localcpus, PMIX_STRING);
20         /* release the list */
21         free(localcpus);
22

```

23 Note that for efficiency, these two values can be computed at the same time.
24 The final *info* array might therefore look like the diagram in 11.5:

25 11.1.4 PMIx_server_deregister_namespace

26 Summary

27 Deregister a namespace.

28 Format

PMIx v1.0

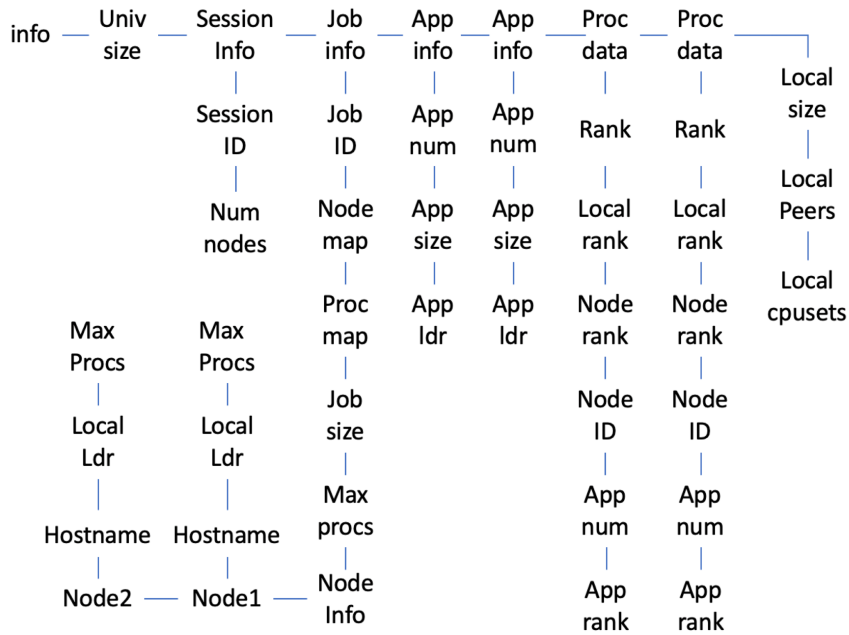


Figure 11.5.: Final information array

```

1  void PMIx_server_deregister_namespace(const pmix_namespace_t namespace,
2                                       pmix_op_cbfunc_t cbfunc, void *cbdata)
3
4  IN  namespace
5       Namespace (string)
6
7  IN  cbfunc
8       Callback function pmix_op_cbfunc_t (function reference)
9
10 IN  cbdata
11      Data to be passed to the callback function (memory reference)
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100

```

1 11.1.5 PMIx_server_register_client

2 Summary

3 Register a client process with the PMIx server library.

4 Format

PMIx v1.0

C

5 `pmix_status_t`

```
6 PMIx_server_register_client(const pmix_proc_t *proc,  
7                            uid_t uid, gid_t gid,  
8                            void *server_object,  
9                            pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

10 **IN** `proc`

11 `pmix_proc_t` structure (handle)

12 **IN** `uid`

13 user id (integer)

14 **IN** `gid`

15 group id (integer)

16 **IN** `server_object`

17 (memory reference)

18 **IN** `cbfunc`

19 Callback function `pmix_op_cbfunc_t` (function reference)

20 **IN** `cbdata`

21 Data to be passed to the callback function (memory reference)

22 Returns one of the following:

- 23 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
24 will be returned in the provided `cbfunc`. Note that the library must not invoke the callback
25 function prior to returning from the API.
- 26 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
27 returned *success* - the `cbfunc` will not be called
- 28 • a PMIx error constant indicating either an error in the input or that the request was immediately
29 processed and failed - the `cbfunc` will not be called

30 Description

31 Register a client process with the PMIx server library.

32 The host server can also, if it desires, provide an object it wishes to be returned when a server
33 function is called that relates to a specific process. For example, the host server may have an object
34 that tracks the specific client. Passing the object to the library allows the library to provide that
35 object to the host server during subsequent calls related to that client, such as a

1 [pmix_server_client_connected_fn_t](#) function. This allows the host server to access
2 the object without performing a lookup based on the client's namespace and rank.

Advice to PMIx server hosts

3 Host environments are required to execute this operation prior to starting the client process. The
4 expected user ID and group ID of the child process allows the server library to properly authenticate
5 clients as they connect by requiring the two values to match. Accordingly, the detected user and
6 group ID's of the connecting process are not included in the
7 [pmix_server_client_connected_fn_t](#) server module function.

Advice to PMIx library implementers

8 For security purposes, the PMIx server library should check the user and group ID's of a
9 connecting process against those provided for the declared client process identifier via the
10 [PMIx_server_register_client](#) prior to completing the connection.

11 11.1.6 [PMIx_server_deregister_client](#)

12 Summary

13 Deregister a client and purge all data relating to it.

14 Format

PMIx v1.0

```
15 void  
16 PMIx_server_deregister_client(const pmix_proc_t *proc,  
17                               pmix_op_cbfunc_t cbfunc, void *cbdata)
```

18 **IN** `proc`
19 [pmix_proc_t](#) structure (handle)
20 **IN** `cbfunc`
21 Callback function [pmix_op_cbfunc_t](#) (function reference)
22 **IN** `cbdata`
23 Data to be passed to the callback function (memory reference)

24 Description

25 The [PMIx_server_deregister_namespace](#) API will delete all client information for that
26 namespace. The PMIx server library will automatically perform that operation upon disconnect of
27 all local clients. This API is therefore intended primarily for use in exception cases, but can be
28 called in non-exception cases if desired. Note that the library must not invoke the callback function
29 prior to returning from the API.

1 11.1.7 PMIx_server_setup_fork

2 Summary

3 Setup the environment of a child process to be forked by the host.

4 Format

PMIx v1.0

C

5 `pmix_status_t`

6 `PMIx_server_setup_fork(const pmix_proc_t *proc,`
7 `char ***env)`

C

8 **IN** `proc`

9 `pmix_proc_t` structure (handle)

10 **IN** `env`

11 Environment array (array of strings)

12 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

13 Description

14 Setup the environment of a child process to be forked by the host so it can correctly interact with
15 the PMIx server.

Advice to PMIx server hosts

16 Host environments are required to execute this operation prior to starting the client process.

17 The PMIx client needs some setup information so it can properly connect back to the server. This
18 function will set appropriate environmental variables for this purpose, and will also provide any
19 environmental variables that were specified in the launch command (e.g., via `PMIx_Spawn`) plus
20 other values (e.g., variables required to properly initialize the client's fabric library).

21 11.1.8 PMIx_server_dmodex_request

22 Summary

23 Define a function by which the host server can request modex data from the local PMIx server.

1
PMIx v1.0

Format

C

```
2 pmix_status_t PMIx_server_dmodex_request(const pmix_proc_t *proc,  
3 pmix_dmodex_response_fn_t cbfunc,  
4 void *cbdata)
```

C

- 5 **IN** `proc`
- 6 `pmix_proc_t` structure (handle)
- 7 **IN** `cbfunc`
- 8 Callback function `pmix_dmodex_response_fn_t` (function reference)
- 9 **IN** `cbdata`
- 10 Data to be passed to the callback function (memory reference)

11 Returns one of the following:

- 12 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
- 13 will be returned in the provided `cbfunc`. Note that the library must not invoke the callback
- 14 function prior to returning from the API.
- 15 • a PMIx error constant indicating an error in the input - the `cbfunc` will not be called

Description

17 Define a function by which the host server can request modex data from the local PMIx server.
18 Traditional wireup procedures revolve around the per-process posting of data (e.g., location and
19 endpoint information) via the **PMIx_Put** and **PMIx_Commit** functions followed by a
20 **PMIx_Fence** barrier that globally exchanges the posted information. However, the barrier
21 operation represents a significant time impact at large scale.

22 PMIx supports an alternative wireup method known as *Direct Modex* that replaces the
23 barrier-based exchange of all process-posted information with on-demand fetch of a peer's data. In
24 place of the barrier operation, data posted by each process is cached on the local PMIx server.
25 When a process requests the information posted by a particular peer, it first checks the local cache
26 to see if the data is already available. If not, then the request is passed to the local PMIx server,
27 which subsequently requests that its RM host request the data from the RM daemon on the node
28 where the specified peer process is located. Upon receiving the request, the RM daemon passes the
29 request into its PMIx server library using the **PMIx_server_dmodex_request** function,
30 receiving the response in the provided `cbfunc` once the indicated process has posted its information.
31 The RM daemon then returns the data to the requesting daemon, who subsequently passes the data
32 to its PMIx server library for transfer to the requesting client.

Advice to users

33 While direct modex allows for faster launch times by eliminating the barrier operation, per-peer
34 retrieval of posted information is less efficient. Optimizations can be implemented - e.g., by
35 returning posted information from all processes on a node upon first request - but in general direct
36 modex remains best suited for sparsely connected applications.

1 11.1.9 PMIx_server_setup_application

2 Summary

3 Provide a function by which the resource manager can request application-specific setup data prior
4 to launch of a `job`.

5 Format

PMIx v2.0

C

```
6 pmix_status_t  
7 PMIx_server_setup_application(const pmix_namespace_t nspace,  
8                             pmix_info_t info[], size_t ninfo,  
9                             pmix_setup_application_cbfunc_t cbfunc,  
10                            void *cbdata)
```

C

11 **IN nspace**

12 namespace (string)

13 **IN info**

14 Array of info structures (array of handles)

15 **IN ninfo**

16 Number of elements in the *info* array (integer)

17 **IN cbfunc**

18 Callback function `pmix_setup_application_cbfunc_t` (function reference)

19 **IN cbdata**

20 Data to be passed to the *cbfunc* callback function (memory reference)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
23 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
24 function prior to returning from the API.
- 25 • a PMIx error constant indicating either an error in the input - the *cbfunc* will not be called

Required Attributes

26 PMIx libraries that support this operation are required to support the following:

27 **PMIX_SETUP_APP_ENVARS** "pmix.setup.env" (bool)

28 Harvest and include relevant environmental variables

29 **PMIX_SETUP_APP_NONENVARS** "pmix.setup.nenv" (bool)

30 Include all relevant data other than environmental variables

31 **PMIX_SETUP_APP_ALL** "pmix.setup.all" (bool)

1 Include all relevant data

2 **PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)

3 Array of `pmix_info_t` describing requested network resources. This must include at
4 least: `PMIX_ALLOC_NETWORK_ID`, `PMIX_ALLOC_NETWORK_TYPE`, and
5 `PMIX_ALLOC_NETWORK_ENDPTS`, plus whatever other descriptors are desired.

6 **PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)

7 The key to be used when accessing this requested network allocation. The allocation will be
8 returned/stored as a `pmix_data_array_t` of `pmix_info_t` indexed by this key and
9 containing at least one entry with the same key and the allocated resource description. The
10 type of the included value depends upon the network support. For example, a TCP allocation
11 might consist of a comma-delimited string of socket ranges such as
12 "32000-32100,33005,38123-38146". Additional entries will consist of any provided
13 resource request directives, along with their assigned values. Examples include:
14 `PMIX_ALLOC_NETWORK_TYPE` - the type of resources provided;
15 `PMIX_ALLOC_NETWORK_PLANE` - if applicable, what plane the resources were assigned
16 from; `PMIX_ALLOC_NETWORK_QOS` - the assigned QoS; `PMIX_ALLOC_BANDWIDTH` -
17 the allocated bandwidth; `PMIX_ALLOC_NETWORK_SEC_KEY` - a security key for the
18 requested network allocation. NOTE: the assigned values may differ from those requested,
19 especially if `PMIX_INFO_REQD` was not set in the request.

20 **PMIX_ALLOC_NETWORK_SEC_KEY** "pmix.alloc.nsec" (`pmix_byte_object_t`)

21 Network security key

22 **PMIX_ALLOC_NETWORK_TYPE** "pmix.alloc.nettype" (char*)

23 Type of desired transport (e.g., "tcp", "udp")

24 **PMIX_ALLOC_NETWORK_PLANE** "pmix.alloc.netplane" (char*)


25 ID string for the NIC (aka *plane*) to be used for this allocation (e.g., CIDR for Ethernet)

26 **PMIX_ALLOC_NETWORK_ENDPTS** "pmix.alloc.endpts" (`size_t`)

27 Number of endpoints to allocate per process

28 **PMIX_ALLOC_NETWORK_ENDPTS_NODE** "pmix.alloc.endpts.nd" (`size_t`)

29 Number of endpoints to allocate per node



30 PMIx libraries that support this operation may support the following:

31 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)

32 Mbits/sec.

33 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)

34 Quality of service level.

35 **PMIX_ALLOC_TIME** "pmix.alloc.time" (`uint32_t`)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29

Time in seconds.

Description

Provide a function by which the RM can request application-specific setup data (e.g., environmental variables, fabric configuration and security credentials) from supporting PMIx server library subsystems prior to initiating launch of a job.

Advice to PMIx server hosts

Host environments are required to execute this operation prior to launching a job. In addition to supported directives, the *info* array must include a description of the **job** using the **PMIX_NODE_MAP** and **PMIX_PROC_MAP** attributes.

This is defined as a non-blocking operation in case contributing subsystems need to perform some potentially time consuming action (e.g., query a remote service) before responding. The returned data must be distributed by the RM and subsequently delivered to the local PMIx server on each node where application processes will execute, prior to initiating execution of those processes.

Advice to PMIx library implementers

Support for harvesting of environmental variables and providing of local configuration information by the PMIx implementation is optional.

11.1.10 PMIx_Register_attributes

Summary

Register host environment attribute support for a function.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Register_attributes(char *function,
                        pmix_regattr_t attrs[],
                        size_t nattrs)
```

C

IN **function**

String name of function (string)

IN **attrs**

Array of **pmix_regattr_t** describing the supported attributes (handle)

IN **nattrs**

Number of elements in *attrs* (**size_t**)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Description

The `PMIx_Register_attributes` function is used by the host environment to register with its PMIx server library the attributes it supports for each `pmix_server_module_t` function. The *function* is the string name of the server module function (e.g., "register_events", "validate_credential", or "allocate") whose attributes are being registered. See the `pmix_regattr_t` entry for a description of the *attrs* array elements.

Note that the host environment can also query the library (using the `PMIx_Query_info_nb` API) for its attribute support both at the server, client, and tool levels once the host has executed `PMIx_server_init` since the server will internally register those values.

Advice to PMIx server hosts

Host environments are strongly encouraged to register all supported attributes immediately after initializing the library to ensure that user requests are correctly serviced.

Advice to PMIx library implementers

PMIx implementations are *required* to register all internally supported attributes for each API during initialization of the library (i.e., when the process calls their respective PMIx init function). Specifically, the implementation *must not* register supported attributes upon first call to a given API as this would prevent users from discovering supported attributes prior to first use of an API.

It is the implementation's responsibility to associate registered attributes for a given `pmix_server_module_t` function with their corresponding user-facing API. Supported attributes *must* be reported to users in terms of their support for user-facing APIs, broken down by the level (see 3.4.33) at which the attribute is supported.

Note that attributes can/will be registered on an API for each level. It is *required* that the implementation support user queries for supported attributes on a per-level basis. Duplicate registrations at the *same* level for a function *shall* return an error - however, duplicate registrations at *different* levels *shall* be independently tracked.

11.1.11 PMIx_server_setup_local_support

Summary

Provide a function by which the local PMIx server can perform any application-specific operations prior to spawning local clients of a given application.

1
PMIx v2.0

Format

C

```

2  pmix_status_t
3  PMIx_server_setup_local_support(const pmix_namespace_t nspace,
4                                pmix_info_t info[], size_t ninfo,
5                                pmix_op_cbfunc_t cbfunc,
6                                void *cbdata);

```

C

- 7 **IN nspace**
Namespace (string)
- 8
- 9 **IN info**
Array of info structures (array of handles)
- 10
- 11 **IN ninfo**
Number of elements in the *info* array (**size_t**)
- 12
- 13 **IN cbfunc**
Callback function [pmix_op_cbfunc_t](#) (function reference)
- 14
- 15 **IN cbdata**
Data to be passed to the callback function (memory reference)
- 16

17 Returns one of the following:

- 18 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
19 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
20 function prior to returning from the API.
- 21 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
22 returned *success* - the *cbfunc* will not be called
- 23 • a PMIx error constant indicating either an error in the input or that the request was immediately
24 processed and failed - the *cbfunc* will not be called

Description

25 Provide a function by which the local PMIx server can perform any application-specific operations
26 prior to spawning local clients of a given application. For example, a network library might need to
27 setup the local driver for “instant on” addressing. The data provided in the *info* array is the data
28 returned to the host RM by the callback function executed as a result of a call to

29 [PMIx_server_setup_application](#).

Advice to PMIx server hosts

31 Host environments are required to execute this operation prior to starting any local application
32 processes from the specified namespace.

1 11.1.12 PMIx_server_IOF_deliver

2 Summary

3 Provide a function by which the host environment can pass forwarded IO to the PMIx server library
4 for distribution to its clients.

5 Format

PMIx v3.0

C

```
6 pmix_status_t  
7 PMIx_server_IOF_deliver(const pmix_proc_t *source,  
8 pmix_iof_channel_t channel,  
9 const pmix_byte_object_t *bo,  
10 const pmix_info_t info[], size_t ninfo,  
11 pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- 12 **IN source**
13 Pointer to [pmix_proc_t](#) identifying source of the IO (handle)
- 14 **IN channel**
15 IO channel of the data ([pmix_iof_channel_t](#))
- 16 **IN bo**
17 Pointer to [pmix_byte_object_t](#) containing the payload to be delivered (handle)
- 18 **IN info**
19 Array of [pmix_info_t](#) metadata describing the data (array of handles)
- 20 **IN ninfo**
21 Number of elements in the *info* array ([size_t](#))
- 22 **IN cbfunc**
23 Callback function [pmix_op_cbfunc_t](#) (function reference)
- 24 **IN cbdata**
25 Data to be passed to the callback function (memory reference)

26 Returns one of the following:

- 27 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
28 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
29 function prior to returning from the API.
- 30 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
31 returned *success* - the *cbfunc* will not be called
- 32 • a PMIx error constant indicating either an error in the input or that the request was immediately
33 processed and failed - the *cbfunc* will not be called

Description

Provide a function by which the host environment can pass forwarded IO to the PMIx server library for distribution to its clients. The PMIx server library is responsible for determining which of its clients have actually registered for the provided data and delivering it. The *cbfunc* callback function will be called once the PMIx server library no longer requires access to the provided data.

11.1.13 PMIx_server_collect_inventory

Summary

Collect inventory of resources on a node

Format

PMIx v3.0

```
pmix_status_t
PMIx_server_collect_inventory(const pmix_info_t directives[],
                             size_t ndirs,
                             pmix_info_cbfunc_t cbfunc,
                             void *cbdata);
```

IN directives

Array of `pmix_info_t` directing the request (array of handles)

IN ndirs

Number of elements in the *directives* array (`size_t`)

IN cbfunc

Callback function to return collected data (`pmix_info_cbfunc_t` function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant. In the event the function returns an error, the *cbfunc* will not be called.

Description

Provide a function by which the host environment can request its PMIx server library collect an inventory of local resources. Supported resources depends upon the PMIx implementation, but may include the local node topology and network interfaces.

Advice to PMIx server hosts

This is a non-blocking API as it may involve somewhat lengthy operations to obtain the requested information. Inventory collection is expected to be a rare event – at system startup and upon command from a system administrator. Inventory updates are expected to initiate a smaller operation involving only the changed information. For example, replacement of a node would generate an event to notify the scheduler with an inventory update without invoking a global inventory operation.

1 11.1.14 PMIx_server_deliver_inventory

2 Summary

3 Pass collected inventory to the PMIx server library for storage

4 Format

PMIx v3.0

C

```
5 pmix_status_t
6 PMIx_server_deliver_inventory(const pmix_info_t info[],
7                               size_t ninfo,
8                               const pmix_info_t directives[],
9                               size_t ndirs,
10                              pmix_op_cbfunc_t cbfunc,
11                              void *cbdata);
```

C

12 **IN info**
13 Array of [pmix_info_t](#) containing the inventory (array of handles)

14 **IN ninfo**
15 Number of elements in the *info* array (**size_t**)

16 **IN directives**
17 Array of [pmix_info_t](#) directing the request (array of handles)

18 **IN ndirs**
19 Number of elements in the *directives* array (**size_t**)

20 **IN cbfunc**
21 Callback function [pmix_op_cbfunc_t](#) (function reference)

22 **IN cbdata**
23 Data to be passed to the callback function (memory reference)

24 Returns one of the following:

- 25 • [PMIX_SUCCESS](#) , indicating that the request is being processed by the host environment - result
26 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
27 function prior to returning from the API.
- 28 • [PMIX_OPERATION_SUCCEEDED](#) , indicating that the request was immediately processed and
29 returned *success* - the *cbfunc* will not be called
- 30 • a PMIx error constant indicating either an error in the input or that the request was immediately
31 processed and failed - the *cbfunc* will not be called

Description

Provide a function by which the host environment can pass inventory information obtained from a node to the PMIx server library for storage. Inventory data is subsequently used by the PMIx server library for allocations in response to `PMIx_server_setup_application`, and may be available to the library's host via the `PMIx_Get` API (depending upon PMIx implementation). The `cbfunc` callback function will be called once the PMIx server library no longer requires access to the provided data.

11.2 Server Function Pointers

PMIx utilizes a "function-shipping" approach to support for implementing the server-side of the protocol. This method allows RMs to implement the server without being burdened with PMIx internal details. When a request is received from the client, the corresponding server function will be called with the information.

Any functions not supported by the RM can be indicated by a `NULL` for the function pointer. Client calls to such functions will return a `PMIX_ERR_NOT_SUPPORTED` status.

The host RM will provide the function pointers in a `pmix_server_module_t` structure passed to `PMIx_server_init`. That module structure and associated function references are defined in this section.

Advice to PMIx server hosts

For performance purposes, the host server is required to return as quickly as possible from all functions. Execution of the function is thus to be done asynchronously so as to allow the PMIx server support library to handle multiple client requests as quickly and scalably as possible.

All data passed to the host server functions is "owned" by the PMIX server support library and must not be free'd. Data returned by the host server via callback function is owned by the host server, which is free to release it upon return from the callback

11.2.1 `pmix_server_module_t` Module

Summary

List of function pointers that a PMIx server passes to `PMIx_server_init` during startup.

Format

C

```
1 typedef struct pmix_server_module_3_0_0_t
2     /* v1x interfaces */
3     pmix_server_client_connected_fn_t    client_connected;
4     pmix_server_client_finalized_fn_t    client_finalized;
5     pmix_server_abort_fn_t               abort;
6     pmix_server_fence_nb_fn_t            fence_nb;
7     pmix_server_dmodex_req_fn_t          direct_modex;
8     pmix_server_publish_fn_t             publish;
9     pmix_server_lookup_fn_t              lookup;
10    pmix_server_unpublish_fn_t            unpublish;
11    pmix_server_spawn_fn_t                spawn;
12    pmix_server_connect_fn_t              connect;
13    pmix_server_disconnect_fn_t           disconnect;
14    pmix_server_register_events_fn_t      register_events;
15    pmix_server_deregister_events_fn_t    deregister_events;
16    pmix_server_listener_fn_t             listener;
17    /* v2x interfaces */
18    pmix_server_notify_event_fn_t         notify_event;
19    pmix_server_query_fn_t                 query;
20    pmix_server_tool_connection_fn_t      tool_connected;
21    pmix_server_log_fn_t                   log;
22    pmix_server_alloc_fn_t                 allocate;
23    pmix_server_job_control_fn_t           job_control;
24    pmix_server_monitor_fn_t               monitor;
25    /* v3x interfaces */
26    pmix_server_get_cred_fn_t              get_credential;
27    pmix_server_validate_cred_fn_t         validate_credential;
28    pmix_server_iof_fn_t                   iof_pull;
29    pmix_server_stdin_fn_t                 push_stdin;
30    pmix_server_module_t;
```

C

32 11.2.2 pmix_server_client_connected_fn_t

33 Summary

34 Notify the host server that a client connected to this server.

1
PMIx v1.0

Format

C

```
typedef pmix_status_t (*pmix_server_client_connected_fn_t) (  
    const pmix_proc_t *proc,  
    void* server_object,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata)
```

C

7 **IN** `proc`
8 `pmix_proc_t` structure (handle)
9 **IN** `server_object`
10 object reference (memory reference)
11 **IN** `cbfunc`
12 Callback function `pmix_op_cbfunc_t` (function reference)
13 **IN** `cbdata`
14 Data to be passed to the callback function (memory reference)

15 Returns one of the following:

- 16 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
17 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
18 prior to returning from the API.
- 19 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
20 returned `success` - the `cbfunc` will not be called
- 21 • a PMIx error constant indicating either an error in the input or that the request was immediately
22 processed and failed - the `cbfunc` will not be called

Description

23 Notify the host environment that a client has called **PMIx_Init**. Note that the client will be in a
24 blocked state until the host server executes the callback function, thus allowing the PMIx server
25 support library to release the client. The `server_object` parameter will be the value of the
26 `server_object` parameter passed to **PMIx_server_register_client** by the host server
27 when registering the connecting client. If provided, an implementation of
28 `pmix_server_client_connected_fn_t` is only required to call the callback function
29 designated. A host server can choose to not be notified when clients connect by setting
30 `pmix_server_client_connected_fn_t` to **NULL**.

31
32 It is possible that only a subset of the clients in a namespace call **PMIx_Init**. The server's
33 `pmix_server_client_connected_fn_t` implementation should not depend on being
34 called once per rank in a namespace or delay calling the callback function until all ranks have
35 connected. However, if a rank makes any PMIx calls, it must first call **PMIx_Init** and therefore
36 the server's `pmix_server_client_connected_fn_t` will be called before any other
37 server functions specific to the rank.

Advice to PMIx server hosts

This operation is an opportunity for a host environment to update the status of the ranks it manages. It is also a convenient and well defined time to perform initialization necessary to support further calls into the server related to that rank.

11.2.3 pmix_server_client_finalized_fn_t

Summary

Notify the host environment that a client called `PMIx_Finalize`.

Format

PMIx v1.0

```
typedef pmix_status_t (*pmix_server_client_finalized_fn_t) (  
    const pmix_proc_t *proc,  
    void* server_object,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata)
```

IN `proc`

`pmix_proc_t` structure (handle)

IN `server_object`

object reference (memory reference)

IN `cbfunc`

Callback function `pmix_op_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- `PMIX_SUCCESS`, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned `success` - the `cbfunc` will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Description

Notify the host environment that a client called `PMIx_Finalize`. Note that the client will be in a blocked state until the host server executes the callback function, thus allowing the PMIx server support library to release the client. The `server_object` parameter will be the value of the `server_object` parameter passed to `PMIx_server_register_client` by the host server when registering the connecting client. If provided, an implementation of `pmix_server_client_finalized_fn_t` is only required to call the callback function designated. A host server can choose to not be notified when clients finalize by setting `pmix_server_client_finalized_fn_t` to `NULL`.

Note that the host server is only being informed that the client has called `PMIx_Finalize`. The client might not have exited. If a client exits without calling `PMIx_Finalize`, the server support library will not call the `pmix_server_client_finalized_fn_t` implementation.

Advice to PMIx server hosts

This operation is an opportunity for a host server to update the status of the tasks it manages. It is also a convenient and well defined time to release resources used to support that client.

11.2.4 `pmix_server_abort_fn_t`

Summary

Notify the host environment that a local client called `PMIx_Abort`.

Format

PMIx v1.0

C

```
typedef pmix_status_t (*pmix_server_abort_fn_t) (  
    const pmix_proc_t *proc,  
    void *server_object,  
    int status,  
    const char msg[],  
    pmix_proc_t procs[],  
    size_t nprocs,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata)
```

1 **IN** **proc**
 2 **pmix_proc_t** structure identifying the process requesting the abort (handle)
 3 **IN** **server_object**
 4 object reference (memory reference)
 5 **IN** **status**
 6 exit status (integer)
 7 **IN** **msg**
 8 exit status message (string)
 9 **IN** **procs**
 10 Array of **pmix_proc_t** structures identifying the processes to be terminated (array of
 11 handles)
 12 **IN** **nprocs**
 13 Number of elements in the *procs* array (integer)
 14 **IN** **cbfunc**
 15 Callback function **pmix_op_cbfunc_t** (function reference)
 16 **IN** **cbdata**
 17 Data to be passed to the callback function (memory reference)

18 Returns one of the following:

- 19 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
 20 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
 21 prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
 23 returned *success* - the *cbfunc* will not be called
- 24 • a PMIx error constant indicating either an error in the input or that the request was immediately
 25 processed and failed - the *cbfunc* will not be called

26 **Description**

27 A local client called **PMIx_Abort** . Note that the client will be in a blocked state until the host
 28 server executes the callback function, thus allowing the PMIx server library to release the client.
 29 The array of *procs* indicates which processes are to be terminated. A **NULL** indicates that all
 30 processes in the client's namespace are to be terminated.

31 **11.2.5 pmix_server_fence_nb_fn_t**

32 **Summary**

33 At least one client called either **PMIx_Fence** or **PMIx_Fence_nb** .

1
PMIx v1.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_fence_fn_t) (  
3     const pmix_proc_t procs[],  
4     size_t nprocs,  
5     const pmix_info_t info[],  
6     size_t ninfo,  
7     char *data, size_t ndata,  
8     pmix_modex_cbfunc_t cbfunc,  
9     void *cbdata)
```

C

10 **IN** `procs`
11 Array of `pmix_proc_t` structures identifying operation participants(array of handles)

12 **IN** `nprocs`
13 Number of elements in the `procs` array (integer)

14 **IN** `info`
15 Array of info structures (array of handles)

16 **IN** `ninfo`
17 Number of elements in the `info` array (integer)

18 **IN** `data`
19 (string)

20 **IN** `ndata`
21 (integer)

22 **IN** `cbfunc`
23 Callback function `pmix_modex_cbfunc_t` (function reference)

24 **IN** `cbdata`
25 Data to be passed to the callback function (memory reference)

26 Returns one of the following:

- 27
- 28 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
29 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
prior to returning from the API.
 - 30 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
31 returned `success` - the `cbfunc` will not be called
 - 32 • a PMIx error constant indicating either an error in the input or that the request was immediately
33 processed and failed - the `cbfunc` will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing.

The following attributes are required to be supported by all host environments:

PMIX_COLLECT_DATA "pmix.collect" (bool)

Collect data and return it at the end of the operation.

Optional Attributes

The following attributes are optional for host environments:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

PMIX_COLLECTIVE_ALGO "pmix.calgo" (char*)

Comma-delimited list of algorithms to use for the collective operation. PMIx does not impose any requirements on a host environment’s collective algorithms. Thus, the acceptable values for this attribute will be environment-dependent - users are encouraged to check their host environment for supported values.

PMIX_COLLECTIVE_ALGO_REQD "pmix.calreqd" (bool)

If **true**, indicates that the requested choice of algorithm is mandatory.

Advice to PMIx server hosts

Host environment are required to return **PMIX_ERR_NOT_SUPPORTED** if passed an attributed marked as **PMIX_INFO_REQD** that they do not support, even if support for that attribute is optional.

Description

All local clients in the provided array of *procs* called either `PMIx_Fence` or `PMIx_Fence_nb`. In either case, the host server will be called via a non-blocking function to execute the specified operation once all participating local processes have contributed. All processes in the specified *procs* array are required to participate in the `PMIx_Fence` / `PMIx_Fence_nb` operation. The callback is to be executed once every daemon hosting at least one participant has called the host server's `pmix_server_fence_nb_fn_t` function.

Advice to PMIx library implementers

The PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective. Data received from each node must be simply concatenated to form an aggregated unit, as shown in the following example:

```
C
uint8_t *blob1, *blob2, *total;
size_t sz_blob1, sz_blob2, sz_total;

sz_total = sz_blob1 + sz_blob2;
total = (uint8_t*)malloc(sz_total);
memcpy(total, blob1, sz_blob1);
memcpy(&total[sz_blob1], blob2, sz_blob2);
```

Note that the ordering of the data blobs does not matter.

The provided data is to be collectively shared with all PMIx servers involved in the fence operation, and returned in the modex *cbfunc*. A **NULL** data value indicates that the local processes had no data to contribute.

The array of *info* structs is used to pass user-requested options to the server. This can include directives as to the algorithm to be used to execute the fence operation. The directives are optional unless the `PMIX_INFO_REQD` flag has been set - in such cases, the host RM is required to return an error if the directive cannot be met.

1 11.2.6 pmix_server_dmodex_req_fn_t

2 Summary

3 Used by the PMIx server to request its local host contact the PMIx server on the remote node that
4 hosts the specified proc to obtain and return a direct modex blob for that proc.

5 Format

PMIx v1.0

C

```
6 typedef pmix_status_t (*pmix_server_dmodex_req_fn_t) (  
7     const pmix_proc_t *proc,  
8     const pmix_info_t info[],  
9     size_t ninfo,  
10    pmix_modex_cbfunc_t cbfunc,  
11    void *cbdata)
```

C

12 **IN** `proc`

13 `pmix_proc_t` structure identifying the process whose data is being requested (handle)

14 **IN** `info`

15 Array of info structures (array of handles)

16 **IN** `ninfo`

17 Number of elements in the `info` array (integer)

18 **IN** `cbfunc`

19 Callback function `pmix_modex_cbfunc_t` (function reference)

20 **IN** `cbdata`

21 Data to be passed to the callback function (memory reference)

22 Returns one of the following:

- 23 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
24 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
25 prior to returning from the API.
- 26 • a PMIx error constant indicating either an error in the input or that the request was immediately
27 processed and failed - the `cbfunc` will not be called

Required Attributes

28 PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Description

Used by the PMIx server to request its local host contact the PMIx server on the remote node that hosts the specified proc to obtain and return any information that process posted via calls to **PMIx_Put** and **PMIx_Commit**.

The array of *info* structs is used to pass user-requested options to the server. This can include a timeout to preclude an indefinite wait for data that may never become available. The directives are optional unless the *mandatory* flag has been set - in such cases, the host RM is required to return an error if the directive cannot be met.

11.2.7 pmix_server_publish_fn_t

Summary

Publish data per the PMIx API specification.

Format

PMIx v1.0

```
typedef pmix_status_t (*pmix_server_publish_fn_t) (  
    const pmix_proc_t *proc,  
    const pmix_info_t info[],  
    size_t ninfo,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata)
```

IN **proc**
pmix_proc_t structure of the process publishing the data (handle)

IN **info**
Array of info structures (array of handles)

IN **ninfo**
Number of elements in the *info* array (integer)

IN **cbfunc**
Callback function **pmix_op_cbfunc_t** (function reference)

1 **IN cldata**

2 Data to be passed to the callback function (memory reference)

3 Returns one of the following:

- 4 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
5 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
6 prior to returning from the API.
- 7 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
8 returned *success* - the *cbfunc* will not be called
- 9 • a PMIx error constant indicating either an error in the input or that the request was immediately
10 processed and failed - the *cbfunc* will not be called

▼----- Required Attributes -----▼

11 PMIx libraries are required to pass any provided attributes to the host environment for processing.
12 In addition, the following attributes are required to be included in the passed *info* array:

13 **PMIX_USERID** "pmix.euid" (uint32_t)
14 Effective user id.

15 **PMIX_GRPID** "pmix.egid" (uint32_t)
16 Effective group id.

17
18 Host environments that implement this entry point are required to support the following attributes:

19 **PMIX_RANGE** "pmix.range" (pmix_data_range_t)
20 Value for calls to publish/lookup/unpublish or for monitoring event notifications.

21 **PMIX_PERSISTENCE** "pmix.persist" (pmix_persistence_t)
22 Value for calls to **PMIx_Publish** .

▲----- Optional Attributes -----▲

23 The following attributes are optional for host environments that support this operation:

24 **PMIX_TIMEOUT** "pmix.timeout" (int)
25 Time in seconds before the specified operation should time out (0 indicating infinite) in
26 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
27 the target process from ever exposing its data.

Description

Publish data per the **PMIx_Publish** specification. The callback is to be executed upon completion of the operation. The default data range is left to the host environment, but expected to be **PMIX_RANGE_SESSION**, and the default persistence **PMIX_PERSIST_SESSION** or their equivalent. These values can be specified by including the respective attributed in the *info* array.

The persistence indicates how long the server should retain the data.

Advice to PMIx server hosts

The host environment is not required to guarantee support for any specific range - i.e., the environment does not need to return an error if the data store doesn't support a specified range so long as it is covered by some internally defined range. However, the server must return an error (a) if the key is duplicative within the storage range, and (b) if the server does not allow overwriting of published info by the original publisher - it is left to the discretion of the host environment to allow info-key-based flags to modify this behavior.

The **PMIX_USERID** and **PMIX_GRPID** of the publishing process will be provided to support authorization-based access to published information and must be returned on any subsequent lookup request.

11.2.8 pmix_server_lookup_fn_t

Summary

Lookup published data.

Format

PMIx v1.0

C

```
typedef pmix_status_t (*pmix_server_lookup_fn_t) (
    const pmix_proc_t *proc,
    char **keys,
    const pmix_info_t info[],
    size_t ninfo,
    pmix_lookup_cbfnc_t cbfunc,
    void *cbdata)
```

1 **IN** **proc**
 2 **pmix_proc_t** structure of the process seeking the data (handle)
 3 **IN** **keys**
 4 (array of strings)
 5 **IN** **info**
 6 Array of info structures (array of handles)
 7 **IN** **ninfo**
 8 Number of elements in the *info* array (integer)
 9 **IN** **cbfunc**
 10 Callback function **pmix_lookup_cbfunc_t** (function reference)
 11 **IN** **cbdata**
 12 Data to be passed to the callback function (memory reference)

13 Returns one of the following:

- 14 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
 15 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
 16 prior to returning from the API.
- 17 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
 18 returned *success* - the *cbfunc* will not be called
- 19 • a PMIx error constant indicating either an error in the input or that the request was immediately
 20 processed and failed - the *cbfunc* will not be called

Required Attributes

21 PMIx libraries are required to pass any provided attributes to the host environment for processing.
 22 In addition, the following attributes are required to be included in the passed *info* array:

23 **PMIX_USERID** "pmix.euid" (**uint32_t**)
 24 Effective user id.

25 **PMIX_GRPID** "pmix.egid" (**uint32_t**)
 26 Effective group id.

27

28 Host environments that implement this entry point are required to support the following attributes:

29 **PMIX_RANGE** "pmix.range" (**pmix_data_range_t**)
 30 Value for calls to publish/lookup/unpublish or for monitoring event notifications.

31 **PMIX_WAIT** "pmix.wait" (**int**)
 32 Caller requests that the PMIx server wait until at least the specified number of values are
 33 found (*0* indicates all and is the default).

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Description

Lookup published data. The host server will be passed a **NULL**-terminated array of string keys identifying the data being requested.

The array of *info* structs is used to pass user-requested options to the server. The default data range is left to the host environment, but expected to be **PMIX_RANGE_SESSION**. This can include a wait flag to indicate that the server should wait for all data to become available before executing the callback function, or should immediately callback with whatever data is available. In addition, a timeout can be specified on the wait to preclude an indefinite wait for data that may never be published.

Advice to PMIx server hosts

The **PMIX_USERID** and **PMIX_GRPID** of the requesting process will be provided to support authorization-based access to published information. The host environment is not required to guarantee support for any specific range - i.e., the environment does not need to return an error if the data store doesn't support a specified range so long as it is covered by some internally defined range.

11.2.9 pmix_server_unpublish_fn_t

Summary

Delete data from the data store.

1
PMIx v1.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_unpublish_fn_t) (  
3     const pmix_proc_t *proc,  
4     char **keys,  
5     const pmix_info_t info[],  
6     size_t ninfo,  
7     pmix_op_cbfunc_t cbfunc,  
8     void *cbdata)
```

C

9 **IN** `proc`
10 `pmix_proc_t` structure identifying the process making the request (handle)
11 **IN** `keys`
12 (array of strings)
13 **IN** `info`
14 Array of info structures (array of handles)
15 **IN** `ninfo`
16 Number of elements in the *info* array (integer)
17 **IN** `cbfunc`
18 Callback function `pmix_op_cbfunc_t` (function reference)
19 **IN** `cbdata`
20 Data to be passed to the callback function (memory reference)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
23 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
24 prior to returning from the API.
- 25 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
26 returned *success* - the *cbfunc* will not be called
- 27 • a PMIx error constant indicating either an error in the input or that the request was immediately
28 processed and failed - the *cbfunc* will not be called

Required Attributes

29 PMIx libraries are required to pass any provided attributes to the host environment for processing.
30 In addition, the following attributes are required to be included in the passed *info* array:

31 **PMIX_USERID** "pmix.euid" (uint32_t)
32 Effective user id.
33 **PMIX_GRPID** "pmix.egid" (uint32_t)
34 Effective group id.

Host environments that implement this entry point are required to support the following attributes:

PMIX_RANGE "pmix.range" (`pmix_data_range_t`)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (`int`)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Description

Delete data from the data store. The host server will be passed a **NULL**-terminated array of string keys, plus potential directives such as the data range within which the keys should be deleted. The default data range is left to the host environment, but expected to be `PMIX_RANGE_SESSION`. The callback is to be executed upon completion of the delete procedure.

Advice to PMIx server hosts

The `PMIX_USERID` and `PMIX_GRPID` of the requesting process will be provided to support authorization-based access to published information. The host environment is not required to guarantee support for any specific range - i.e., the environment does not need to return an error if the data store doesn't support a specified range so long as it is covered by some internally defined range.

11.2.10 `pmix_server_spawn_fn_t`

Summary

Spawn a set of applications/processes as per the `PMIx_Spawn` API.

1
PMIx v1.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_spawn_fn_t) (  
3     const pmix_proc_t *proc,  
4     const pmix_info_t job_info[],  
5     size_t ninfo,  
6     const pmix_app_t apps[],  
7     size_t napps,  
8     pmix_spawn_cbfunc_t cbfunc,  
9     void *cbdata)
```

C

10 **IN** `proc`
11 `pmix_proc_t` structure of the process making the request (handle)
12 **IN** `job_info`
13 Array of info structures (array of handles)
14 **IN** `ninfo`
15 Number of elements in the `jobinfo` array (integer)
16 **IN** `apps`
17 Array of `pmix_app_t` structures (array of handles)
18 **IN** `napps`
19 Number of elements in the `apps` array (integer)
20 **IN** `cbfunc`
21 Callback function `pmix_spawn_cbfunc_t` (function reference)
22 **IN** `cbdata`
23 Data to be passed to the callback function (memory reference)

24 Returns one of the following:

- 25 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
26 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
27 prior to returning from the API.
- 28 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
29 returned `success` - the `cbfunc` will not be called
- 30 • a PMIx error constant indicating either an error in the input or that the request was immediately
31 processed and failed - the `cbfunc` will not be called

Required Attributes

32 PMIx libraries are required to pass any provided attributes to the host environment for processing.
33 In addition, the following attributes are required to be included in the passed `info` array:

34 **PMIX_USERID** "`pmix.euid`" (`uint32_t`)
35 Effective user id.

1 **PMIX_GRPID** "pmix.egid" (uint32_t)

2 Effective group id.

3
4 Host environments that provide this module entry point are required to pass the **PMIX_SPAWNED**
5 and **PMIX_PARENT_ID** attributes to all PMIx servers launching new child processes so those
6 values can be returned to clients upon connection to the PMIx server. In addition, they are required
7 to support the following attributes when present in either the *job_info* or the *info* array of an
8 element of the *apps* array:

9 **PMIX_WDIR** "pmix.wdir" (char*)

10 Working directory for spawned processes.

11 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

12 Set the application's current working directory to the session working directory assigned by
13 the RM - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for
14 the rank to discover the session working directory assigned to the provided namespace

15 **PMIX_PREFIX** "pmix.prefix" (char*)

16 Prefix to use for starting spawned processes.

17 **PMIX_HOST** "pmix.host" (char*)

18 Comma-delimited list of hosts to use for spawned processes.

19 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

20 Hostfile to use for spawned processes.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

21 The following attributes are optional for host environments that support this operation:

22 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

23 Hostfile listing hosts to add to existing allocation.

24 **PMIX_ADD_HOST** "pmix.addhost" (char*)

25 Comma-delimited list of hosts to add to the allocation.

26 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

27 Preload binaries onto nodes.

28 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

29 Comma-delimited list of files to pre-position on nodes.

30 **PMIX_PERSONALITY** "pmix.pers" (char*)

31 Name of personality to use.

32 **PMIX_MAPPER** "pmix.mapper" (char*)

1 Mapping mechanism to use for placing spawned processes - when accessed using
2 `PMIx_Get` , use the `PMIX_RANK_WILDCARD` value for the rank to discover the mapping
3 mechanism used for the provided namespace.

4 `PMIX_DISPLAY_MAP` "pmix.dispmap" (bool)

5 Display process mapping upon spawn.

6 `PMIX_PPR` "pmix.ppr" (char*)

7 Number of processes to spawn on each identified resource.

8 `PMIX_MAPBY` "pmix.mapby" (char*)

9 Process mapping policy - when accessed using `PMIx_Get` , use the
10 `PMIX_RANK_WILDCARD` value for the rank to discover the mapping policy used for the
11 provided namespace

12 `PMIX_RANKBY` "pmix.rankby" (char*)

13 Process ranking policy - when accessed using `PMIx_Get` , use the
14 `PMIX_RANK_WILDCARD` value for the rank to discover the ranking algorithm used for the
15 provided namespace

16 `PMIX_BINDTO` "pmix.bindto" (char*)

17 Process binding policy - when accessed using `PMIx_Get` , use the
18 `PMIX_RANK_WILDCARD` value for the rank to discover the binding policy used for the
19 provided namespace

20 `PMIX_NON_PMI` "pmix.nonpmi" (bool)

21 Spawned processes will not call `PMIx_Init` .

22 `PMIX_STDIN_TGT` "pmix.stdin" (uint32_t)

23 Spawned process rank that is to receive `stdin`.

24 `PMIX_FWD_STDIN` "pmix.fwd.stdin" (bool)

25 Forward this process's `stdin` to the designated process.

26 `PMIX_FWD_STDOUT` "pmix.fwd.stdout" (bool)

27 Forward `stdout` from spawned processes to this process.

28 `PMIX_FWD_STDERR` "pmix.fwd.stderr" (bool)

29 Forward `stderr` from spawned processes to this process.

30 `PMIX_DEBUGGER_DAEMONS` "pmix.debugger" (bool)

31 Spawned application consists of debugger daemons.

32 `PMIX_TAG_OUTPUT` "pmix.tagout" (bool)

33 Tag application output with the identity of the source process.

34 `PMIX_TIMESTAMP_OUTPUT` "pmix.tsout" (bool)

35 Timestamp output from applications.

36 `PMIX_MERGE_STDERR_STDOUT` "pmix.mergeerrout" (bool)

1 Merge `stdout` and `stderr` streams from application processes.

2 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)

3 Output application output to the specified file.

4 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)

5 Mark the `argv` with the rank of the process.

6 **PMIX_CPUS_PER_PROC** "pmix.cputerproc" (uint32_t)

7 Number of cpus to assign to each rank - when accessed using `PMIx_Get`, use the
8 **PMIX_RANK_WILDCARD** value for the rank to discover the cpus/process assigned to the
9 provided namespace

10 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)

11 Do not place processes on the head node.

12 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)

13 Do not oversubscribe the cpus.

14 **PMIX_REPORT_BINDINGS** "pmix.repbind" (bool)

15 Report bindings of the individual processes.

16 **PMIX_CPU_LIST** "pmix.cpulist" (char*)

17 List of cpus to use for this job - when accessed using `PMIx_Get`, use the
18 **PMIX_RANK_WILDCARD** value for the rank to discover the cpu list used for the provided
19 namespace

20 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)

21 Application supports recoverable operations.

22 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)

23 Application is continuous, all failed processes should be immediately restarted.

24 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)

25 Maximum number of times to restart a job - when accessed using `PMIx_Get`, use the
26 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
27 namespace

28 **PMIX_TIMEOUT** "pmix.timeout" (int)

29 Time in seconds before the specified operation should time out (0 indicating infinite) in
30 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
31 the target process from ever exposing its data.



Description

Spawn a set of applications/processes as per the [PMIx_Spawn](#) API. Note that applications are not required to be MPI or any other programming model. Thus, the host server cannot make any assumptions as to their required support. The callback function is to be executed once all processes have been started. An error in starting any application or process in this request shall cause all applications and processes in the request to be terminated, and an error returned to the originating caller.

Note that a timeout can be specified in the `job_info` array to indicate that failure to start the requested job within the given time should result in termination to avoid hangs.

11.2.11 pmix_server_connect_fn_t

Summary

Record the specified processes as *connected*.

Format

PMIx v1.0

```
typedef pmix_status_t (*pmix_server_connect_fn_t) (  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t info[],  
    size_t ninfo,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata)
```

IN procs

Array of [pmix_proc_t](#) structures identifying participants (array of handles)

IN nprocs

Number of elements in the *procs* array (integer)

IN info

Array of info structures (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

IN cbfunc

Callback function [pmix_op_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.

- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

The following attributes are optional for host environments that support this operation:

`PMIX_TIMEOUT` "`pmix.timeout`" (`int`)

Time in seconds before the specified operation should time out (`0` indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

`PMIX_COLLECTIVE_ALGO` "`pmix.calgo`" (`char*`)

Comma-delimited list of algorithms to use for the collective operation. PMIx does not impose any requirements on a host environment’s collective algorithms. Thus, the acceptable values for this attribute will be environment-dependent - users are encouraged to check their host environment for supported values.

`PMIX_COLLECTIVE_ALGO_REQD` "`pmix.calreqd`" (`bool`)

If `true`, indicates that the requested choice of algorithm is mandatory.

Description

Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The callback is to be executed once every daemon hosting at least one participant has called the host server’s `pmix_server_connect_fn_t` function, and the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes.

Advice to PMIx library implementers

The PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

1 11.2.12 pmix_server_disconnect_fn_t

2 Summary

3 Disconnect a previously connected set of processes.

4 Format

PMIx v1.0

C

```
5 typedef pmix_status_t (*pmix_server_disconnect_fn_t) (  
6     const pmix_proc_t procs[],  
7     size_t nprocs,  
8     const pmix_info_t info[],  
9     size_t ninfo,  
10    pmix_op_cbfunc_t cbfunc,  
11    void *cbdata)
```

C

12 **IN** `procs`

Array of `pmix_proc_t` structures identifying participants (array of handles)

13 **IN** `nprocs`

Number of elements in the `procs` array (integer)

14 **IN** `info`

Array of info structures (array of handles)

15 **IN** `ninfo`

Number of elements in the `info` array (integer)

16 **IN** `cbfunc`

Callback function `pmix_op_cbfunc_t` (function reference)

17 **IN** `cbdata`

Data to be passed to the callback function (memory reference)

18 Returns one of the following:

- 25 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
26 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
27 prior to returning from the API.
- 28 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
29 returned `success` - the `cbfunc` will not be called
- 30 • a PMIx error constant indicating either an error in the input or that the request was immediately
31 processed and failed - the `cbfunc` will not be called

Required Attributes

32 PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Description

Disconnect a previously connected set of processes. The callback is to be executed once every daemon hosting at least one participant has called the host server’s `pmix_server_disconnect_fn_t` function, and the host environment has completed any required supporting operations.

Advice to PMIx library implementers

The PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

A **PMIX_ERR_INVALID_OPERATION** error must be returned if the specified set of *procs* was not previously *connected* via a call to the `pmix_server_connect_fn_t` function.

11.2.13 pmix_server_register_events_fn_t

Summary

Register to receive notifications for the specified events.

1
PMIx v1.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_register_events_fn_t) (  
3     pmix_status_t *codes,  
4     size_t ncodes,  
5     const pmix_info_t info[],  
6     size_t ninfo,  
7     pmix_op_cbfunc_t cbfunc,  
8     void *cbdata)
```

C

- 9 **IN codes**
Array of [pmix_status_t](#) values (array of handles)
- 10 **IN ncodes**
Number of elements in the *codes* array (integer)
- 11 **IN info**
Array of info structures (array of handles)
- 12 **IN ninfo**
Number of elements in the *info* array (integer)
- 13 **IN cbfunc**
Callback function [pmix_op_cbfunc_t](#) (function reference)
- 14 **IN cbdata**
Data to be passed to the callback function (memory reference)

15 Returns one of the following:

- 16 • [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- 17 • [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- 18 • a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

19 PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

- 20 **PMIX_USERID** "pmix.euid" (uint32_t)
Effective user id.
- 21 **PMIX_GRPID** "pmix.egid" (uint32_t)
Effective group id.

Description

Register to receive notifications for the specified status codes. The *info* array included in this API is reserved for possible future directives to further steer notification.

Advice to PMIx library implementers

The PMIx server library must track all client registrations for subsequent notification. This module function shall only be called when:

- the client has requested notification of an environmental code (i.e., a PMIx code in the range beyond `PMIX_ERR_SYS_OTHER`) or a code that lies outside the defined PMIx range of constants; and
- the PMIx server library has not previously requested notification of that code - i.e., the host environment is to be contacted only once a given unique code value

Advice to PMIx server hosts

The host environment is required to pass to its PMIx server library all non-environmental events that directly relate to a registered namespace without the PMIx server library explicitly requesting them. Environmental events are to be translated to their nearest PMIx equivalent code as defined in the range between `PMIX_ERR_SYS_BASE` and `PMIX_ERR_SYS_OTHER` (inclusive).

11.2.14 `pmix_server_deregister_events_fn_t`

Summary

Deregister to receive notifications for the specified events.

1
PMIx v1.0

Format

C

```

typedef pmix_status_t (*pmix_server_deregister_events_fn_t) (
    pmix_status_t *codes,
    size_t ncodes,
    pmix_op_cbfunc_t cbfunc,
    void *cbdata)

```

C

- 7 **IN codes**
Array of `pmix_status_t` values (array of handles)
- 9 **IN ncodes**
Number of elements in the `codes` array (integer)
- 11 **IN cbfunc**
Callback function `pmix_op_cbfunc_t` (function reference)
- 13 **IN cbdata**
Data to be passed to the callback function (memory reference)

15 Returns one of the following:

- 17 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- 19 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned `success` - the `cbfunc` will not be called
- 21 • a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Description

24 Deregister to receive notifications for the specified events to which the PMIx server has previously
25 registered.

Advice to PMIx library implementers

26 The PMIx server library must track all client registrations. This module function shall only be
27 called when:

- 29 • the library is deregistering environmental codes (i.e., a PMIx codes in the range between
30 **PMIX_ERR_SYS_BASE** and **PMIX_ERR_SYS_OTHER**, inclusive) or codes that lies outside
the defined PMIx range of constants; and
- 31 • no client (including the server library itself) remains registered for notifications on any included
32 code - i.e., a code should be included in this call only when no registered notifications against it
33 remain.

1 11.2.15 pmix_server_notify_event_fn_t

2 Summary

3 Notify the specified processes of an event.

4 Format

PMIx v2.0

C

```
5 typedef pmix_status_t (*pmix_server_notify_event_fn_t) (pmix_status_t code,  
6                                                         const pmix_proc_t *source,  
7                                                         pmix_data_range_t range,  
8                                                         pmix_info_t info[],  
9                                                         size_t ninfo,  
10                                                        pmix_op_cbfunc_t cbfunc,  
11                                                        void *cbdata);
```

C

12 **IN code**

13 The [pmix_status_t](#) event code being referenced structure (handle)

14 **IN source**

15 [pmix_proc_t](#) of process that generated the event (handle)

16 **IN range**

17 [pmix_data_range_t](#) range over which the event is to be distributed (handle)

18 **IN info**

19 Optional array of [pmix_info_t](#) structures containing additional information on the event
20 (array of handles)

21 **IN ninfo**

22 Number of elements in the *info* array (integer)

23 **IN cbfunc**

24 Callback function [pmix_op_cbfunc_t](#) (function reference)

25 **IN cbdata**

26 Data to be passed to the callback function (memory reference)

27 Returns one of the following:

- 28 • [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result
29 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
30 prior to returning from the API.
- 31 • [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and
32 returned *success* - the *cbfunc* will not be called
- 33 • a PMIx error constant indicating either an error in the input or that the request was immediately
34 processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing.

Host environments that provide this module entry point are required to support the following attributes:

PMIX_RANGE "pmix.range" (`pmix_data_range_t`)

Value for calls to publish/lookup/unpublish or for monitoring event notifications.

Description

Notify the specified processes (described through a combination of *range* and attributes provided in the *info* array) of an event generated either by the PMIx server itself or by one of its local clients. The process generating the event is provided in the *source* parameter, and any further descriptive information is included in the *info* array.

Advice to PMIx server hosts

The callback function is to be executed once the host environment no longer requires that the PMIx server library maintain the provided data structures. It does not necessarily indicate that the event has been delivered to any process, nor that the event has been distributed for delivery

11.2.16 `pmix_server_listener_fn_t`

Summary

Register a socket the host server can monitor for connection requests.

Format

PMIx v1.0

```
typedef pmix_status_t (*pmix_server_listener_fn_t) (  
    int listening_sd,  
    pmix_connection_cbfunc_t cbfunc,  
    void *cbdata)
```

IN `incoming_sd`

(integer)

IN `cbfunc`

Callback function `pmix_connection_cbfunc_t` (function reference)

IN `cbdata`

(memory reference)

Returns `PMIX_SUCCESS` indicating that the request is accepted, or a negative value corresponding to a PMIx error constant indicating that the request has been rejected.

Description

Register a socket the host environment can monitor for connection requests, harvest them, and then call the PMIx server library's internal callback function for further processing. A listener thread is essential to efficiently harvesting connection requests from large numbers of local clients such as occur when running on large SMPs. The host server listener is required to call `accept` on the incoming connection request, and then pass the resulting socket to the provided `cbfunc`. A `NULL` for this function will cause the internal PMIx server to spawn its own listener thread.

11.2.17 `pmix_server_query_fn_t`

Summary

Query information from the resource manager.

Format

PMIx v2.0

```
typedef pmix_status_t (*pmix_server_query_fn_t) (  
    pmix_proc_t *proct,  
    pmix_query_t *queries, size_t nqueries,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata)
```

IN `proct`

`pmix_proc_t` structure of the requesting process (handle)

IN `queries`

Array of `pmix_query_t` structures (array of handles)

IN `nqueries`

Number of elements in the `queries` array (integer)

IN `cbfunc`

Callback function `pmix_info_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- `PMIX_SUCCESS`, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned `success` - the `cbfunc` will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user id.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group id.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_QUERY_NAMESPACES "pmix.qry.ns" (char*)

Request a comma-delimited list of active namespaces.

PMIX_QUERY_JOB_STATUS "pmix.qry.jst" (pmix_status_t)

Status of a specified, currently executing job.

PMIX_QUERY_QUEUE_LIST "pmix.qry.qlst" (char*)

Request a comma-delimited list of scheduler queues.

PMIX_QUERY_QUEUE_STATUS "pmix.qry.qst" (TBD)

Status of a specified scheduler queue.

PMIX_QUERY_PROC_TABLE "pmix.qry.phtable" (char*)

Input namespace of the job whose information is being requested returns (**pmix_data_array_t**) an array of **pmix_proc_info_t** .

PMIX_QUERY_LOCAL_PROC_TABLE "pmix.qry.lhtable" (char*)

Input namespace of the job whose information is being requested returns (**pmix_data_array_t**) an array of **pmix_proc_info_t** for processes in job on same node.

PMIX_QUERY_SPAWN_SUPPORT "pmix.qry.spawn" (bool)

Return a comma-delimited list of supported spawn attributes.

PMIX_QUERY_DEBUG_SUPPORT "pmix.qry.debug" (bool)

Return a comma-delimited list of supported debug attributes.

PMIX_QUERY_MEMORY_USAGE "pmix.qry.mem" (bool)

Return information on memory usage for the processes indicated in the qualifiers.

PMIX_QUERY_LOCAL_ONLY "pmix.qry.local" (bool)

Constrain the query to local information only.

PMIX_QUERY_REPORT_AVG "pmix.qry.avg" (bool)

Report only average values for sampled information.

1 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
 2 Report minimum and maximum values.

3 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 4 String identifier of the allocation whose status is being requested.

5 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
 6 Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.
 7



8 **Description**

9 Query information from the host environment. The query will include the namespace/rank of the
 10 process that is requesting the info, an array of **pmix_query_t** describing the request, and a
 11 callback function/data for the return.



12 The PMIx server library should not block in this function as the host environment may, depending
 13 upon the information being requested, require significant time to respond.



14 **11.2.18 pmix_server_tool_connection_fn_t**

15 **Summary**

16 Register that a tool has connected to the server.

17 **Format**

PMIx v2.0

```

18 typedef void (*pmix_server_tool_connection_fn_t) (
19             pmix_info_t info[], size_t ninfo,
20             pmix_tool_connection_cbfnc_t cbfunc,
21             void *cbdata)
  
```

- 22 **IN info**
 23 Array of **pmix_info_t** structures (array of handles)
- 24 **IN ninfo**
 25 Number of elements in the *info* array (integer)
- 26 **IN cbfunc**
 27 Callback function **pmix_tool_connection_cbfnc_t** (function reference)
- 28 **IN cbdata**
 29 Data to be passed to the callback function (memory reference)

Required Attributes

PMIx libraries are required to pass the following attributes in the *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user id.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group id.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_FWD_STDOUT "pmix.fwd.stdout" (bool)

Forward **stdout** from spawned processes to this process.

PMIX_FWD_STDERR "pmix.fwd.stderr" (bool)

Forward **stderr** from spawned processes to this process.

PMIX_FWD_STDIN "pmix.fwd.stdin" (bool)

Forward this process's **stdin** to the designated process.

Description

Register that a tool has connected to the server, and request that the tool be assigned a namespace/rank identifier for further interactions. The `pmix_info_t` array is used to pass qualifiers for the connection request, including the effective uid and gid of the calling tool for authentication purposes.

Advice to PMIx server hosts

The host environment is solely responsible for authenticating and authorizing the connection, and for authorizing all subsequent tool requests. The host must not execute the callback function prior to returning from the API.

11.2.19 pmix_server_log_fn_t

Summary

Log data on behalf of a client.

1
PMIx v2.0

Format

C

```
2 typedef void (*pmix_server_log_fn_t) (  
3     const pmix_proc_t *client,  
4     const pmix_info_t data[], size_t ndata,  
5     const pmix_info_t directives[], size_t ndirs,  
6     pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

7 **IN client**
8 `pmix_proc_t` structure (handle)
9 **IN data**
10 Array of info structures (array of handles)
11 **IN ndata**
12 Number of elements in the *data* array (integer)
13 **IN directives**
14 Array of info structures (array of handles)
15 **IN ndirs**
16 Number of elements in the *directives* array (integer)
17 **IN cbfunc**
18 Callback function `pmix_op_cbfunc_t` (function reference)
19 **IN cbdata**
20 Data to be passed to the callback function (memory reference)

Required Attributes

21 PMIx libraries are required to pass any provided attributes to the host environment for processing.
22 In addition, the following attributes are required to be included in the passed *info* array:

23 **PMIX_USERID** "pmix.euid" (`uint32_t`)
24 Effective user id.

25 **PMIX_GRPID** "pmix.egid" (`uint32_t`)
26 Effective group id.


27
28 Host environments that provide this module entry point are required to support the following
29 attributes:

30 **PMIX_LOG_STDERR** "pmix.log.stderr" (`char*`)
31 Log string to `stderr`.

32 **PMIX_LOG_STDOUT** "pmix.log.stdout" (`char*`)
33 Log string to `stdout`.

34 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (`char*`)

1 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
2 otherwise to local syslog



Optional Attributes

3 The following attributes are optional for host environments that support this operation:

4 **PMIX_LOG_MSG** "pmix.log.msg" (**pmix_byte_object_t**)
5 Message blob to be sent somewhere.

6 **PMIX_LOG_EMAIL** "pmix.log.email" (**pmix_data_array_t**)
7 Log via email based on **pmix_info_t** containing directives.

8 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (**char***)
9 Comma-delimited list of email addresses that are to receive the message.

10 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (**char***)
11 Subject line for email.

12 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (**char***)
13 Message to be included in email.



14 **Description**

15 Log data on behalf of a client. This function is not intended for output of computational results, but
16 rather for reporting status and error messages. The host must not execute the callback function prior
17 to returning from the API.

18 **11.2.20 pmix_server_alloc_fn_t**

19 **Summary**

20 Request allocation operations on behalf of a client.

1
PMIx v2.0

Format

C

```

2 typedef pmix_status_t (*pmix_server_alloc_fn_t) (
3     const pmix_proc_t *client,
4     pmix_alloc_directive_t directive,
5     const pmix_info_t data[], size_t ndata,
6     pmix_info_cbfunc_t cbfunc, void *cbdata)

```

C

- 7 **IN client**
- 8 `pmix_proc_t` structure of process making request (handle)
- 9 **IN directive**
- 10 Specific action being requested (`pmix_alloc_directive_t`)
- 11 **IN data**
- 12 Array of info structures (array of handles)
- 13 **IN ndata**
- 14 Number of elements in the *data* array (integer)
- 15 **IN cbfunc**
- 16 Callback function `pmix_info_cbfunc_t` (function reference)
- 17 **IN cbdata**
- 18 Data to be passed to the callback function (memory reference)

19 Returns one of the following:

- 20 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
- 21 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
- 22 prior to returning from the API.
- 23 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
- 24 returned *success* - the *cbfunc* will not be called
- 25 • a PMIx error constant indicating either an error in the input or that the request was immediately
- 26 processed and failed - the *cbfunc* will not be called

Required Attributes

27 PMIx libraries are required to pass any provided attributes to the host environment for processing.
28 In addition, the following attributes are required to be included in the passed *info* array:

- 29 **PMIX_USERID** "pmix.euid" (`uint32_t`)
- 30 Effective user id.
- 31 **PMIX_GRPID** "pmix.egid" (`uint32_t`)
- 32 Effective group id.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35

Host environments that provide this module entry point are required to support the following attributes:

- PMIX_ALLOC_ID** "pmix.alloc.id" (char*)
Provide a string identifier for this allocation request which can later be used to query status of the request.
- PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)
The number of nodes.
- PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
Number of cpus.
- PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
Time in seconds.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

The following attributes are optional for host environments that support this operation:

- PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
Regular expression of the specific nodes.
- PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
Regular expression of the number of cpus for each node.
- PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
Regular expression of the specific cpus indicating the cpus involved.
- PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
Number of Megabytes.
- PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)
Array of `pmix_info_t` describing requested network resources. This must include at least: **PMIX_ALLOC_NETWORK_ID**, **PMIX_ALLOC_NETWORK_TYPE**, and **PMIX_ALLOC_NETWORK_ENDPTS**, plus whatever other descriptors are desired.
- PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)
The key to be used when accessing this requested network allocation. The allocation will be returned/stored as a `pmix_data_array_t` of `pmix_info_t` indexed by this key and containing at least one entry with the same key and the allocated resource description. The type of the included value depends upon the network support. For example, a TCP allocation might consist of a comma-delimited string of socket ranges such as "32000-32100,33005,38123-38146". Additional entries will consist of any provided resource request directives, along with their assigned values. Examples include:
 - PMIX_ALLOC_NETWORK_TYPE** - the type of resources provided;
 - PMIX_ALLOC_NETWORK_PLANE** - if applicable, what plane the resources were assigned

1 from; `PMIX_ALLOC_NETWORK_QOS` - the assigned QoS; `PMIX_ALLOC_BANDWIDTH` -
2 the allocated bandwidth; `PMIX_ALLOC_NETWORK_SEC_KEY` - a security key for the
3 requested network allocation. NOTE: the assigned values may differ from those requested,
4 especially if `PMIX_INFO_REQD` was not set in the request.

5 `PMIX_ALLOC_BANDWIDTH` "pmix.alloc.bw" (float)
6 Mbits/sec.

7 `PMIX_ALLOC_NETWORK_QOS` "pmix.alloc.netqos" (char*)
8 Quality of service level.



9 Description

10 Request new allocation or modifications to an existing allocation on behalf of a client. Several
11 broad categories are envisioned, including the ability to:

- 12 • Request allocation of additional resources, including memory, bandwidth, and compute for an
13 existing allocation. Any additional allocated resources will be considered as part of the current
14 allocation, and thus will be released at the same time.
- 15 • Request a new allocation of resources. Note that the new allocation will be disjoint from (i.e., not
16 affiliated with) the allocation of the requestor - thus the termination of one allocation will not
17 impact the other.
- 18 • Extend the reservation on currently allocated resources, subject to scheduling availability and
19 priorities.
- 20 • Return no-longer-required resources to the scheduler. This includes the *loan* of resources back to
21 the scheduler with a promise to return them upon subsequent request.

22 The callback function provides a *status* to indicate whether or not the request was granted, and to
23 provide some information as to the reason for any denial in the `pmix_info_cbfunc_t` array of
24 `pmix_info_t` structures.

25 11.2.21 pmix_server_job_control_fn_t

26 Summary

27 Execute a job control action on behalf of a client.

1
PMIx v2.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_job_control_fn_t) (  
3     const pmix_proc_t *requestor,  
4     const pmix_proc_t targets[], size_t ntargets,  
5     const pmix_info_t directives[], size_t ndirs,  
6     pmix_info_cbfunc_t cbfunc, void *cbdata)
```

C

7 **IN requestor**
8 [pmix_proc_t](#) structure of requesting process (handle)
9 **IN targets**
10 Array of proc structures (array of handles)
11 **IN ntargets**
12 Number of elements in the *targets* array (integer)
13 **IN directives**
14 Array of info structures (array of handles)
15 **IN ndirs**
16 Number of elements in the *info* array (integer)
17 **IN cbfunc**
18 Callback function [pmix_op_cbfunc_t](#) (function reference)
19 **IN cbdata**
20 Data to be passed to the callback function (memory reference)

21 Returns one of the following:

- 22 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
23 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
24 prior to returning from the API.
- 25 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
26 returned *success* - the *cbfunc* will not be called
- 27 • a PMIx error constant indicating either an error in the input or that the request was immediately
28 processed and failed - the *cbfunc* will not be called

Required Attributes

29 PMIx libraries are required to pass any attributes provided by the client to the host environment for
30 processing. In addition, the following attributes are required to be included in the passed *info* array:

31 **PMIX_USERID** "pmix.euid" (uint32_t)
32 Effective user id.
33 **PMIX_GRPID** "pmix.egid" (uint32_t)
34 Effective group id.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

Host environments that provide this module entry point are required to support the following attributes:

- PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)
Provide a string identifier for this request. The user can provide an identifier for the requested operation, thus allowing them to later request status of the operation or to terminate it. The host, therefore, shall track it with the request for future reference.
- PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)
Pause the specified processes.
- PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)
Resume ("un-pause") the specified processes.
- PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
Forcibly terminate the specified processes and cleanup.
- PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
Send given signal to specified processes.
- PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
Politely terminate the specified processes.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

The following attributes are optional for host environments that support this operation:

- PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
Cancel the specified request - the provided request ID must match the **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control** . An ID of **NULL** implies cancel all requests from this requestor.
- PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
Restart the specified processes using the given checkpoint ID.
- PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
Checkpoint the specified processes and assign the given ID to it.
- PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
Use event notification to trigger a process checkpoint.
- PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
Use the given signal to trigger a process checkpoint.
- PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
Time in seconds to wait for a checkpoint to complete.

1 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
 2 "pmix.jctrl.ckmethod" (pmix_data_array_t)
 3 Array of pmix_info_t declaring each method and value supported by this application.
 4 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
 5 Regular expression identifying nodes that are to be provisioned.
 6 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
 7 Name of the image that is to be provisioned.
 8 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
 9 Indicate that the job can be pre-empted.



10 **Description**

11 Execute a job control action on behalf of a client. The *targets* array identifies the processes to
 12 which the requested job control action is to be applied. A **NULL** value can be used to indicate all
 13 processes in the caller's namespace. The use of **PMIX_RANK_WILDCARD** can also be used to
 14 indicate that all processes in the given namespace are to be included.

15 The directives are provided as pmix_info_t structures in the *directives* array. The callback
 16 function provides a *status* to indicate whether or not the request was granted, and to provide some
 17 information as to the reason for any denial in the pmix_info_cbfunc_t array of
 18 pmix_info_t structures.

19 **11.2.22 pmix_server_monitor_fn_t**

20 **Summary**

21 Request that a client be monitored for activity.

22 **Format**

PMIx v2.0

23 **typedef pmix_status_t (*pmix_server_monitor_fn_t) (**
 24 **const pmix_proc_t *requestor,**
 25 **const pmix_info_t *monitor, pmix_status_t error**
 26 **const pmix_info_t directives[], size_t ndirs,**
 27 **pmix_info_cbfunc_t cbfunc, void *cbdata);**

28 **IN requestor**

pmix_proc_t structure of requesting process (handle)

30 **IN monitor**

pmix_info_t identifying the type of monitor being requested (handle)

32 **IN error**

Status code to use in generating event if alarm triggers (integer)

1 **IN directives**
2 Array of info structures (array of handles)
3 **IN ndirs**
4 Number of elements in the *info* array (integer)
5 **IN cbfunc**
6 Callback function `pmix_op_cbfunc_t` (function reference)
7 **IN cbdata**
8 Data to be passed to the callback function (memory reference)

9 Returns one of the following:

- 10 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
11 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
12 prior to returning from the API.
- 13 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
14 returned *success* - the *cbfunc* will not be called
- 15 • a PMIx error constant indicating either an error in the input or that the request was immediately
16 processed and failed - the *cbfunc* will not be called

17 This entry point is only called for monitoring requests that are not directly supported by the PMIx
18 server library itself.

▼----- Required Attributes -----▼

19 If supported by the PMIx server library, then the library must not pass any supported attributes to
20 the host environment. Any attributes provided by the client that are not directly supported by the
21 server library must be passed to the host environment if it provides this module entry. In addition,
22 the following attributes are required to be included in the passed *info* array:

23 **PMIX_USERID** "pmix.euid" (uint32_t)
24 Effective user id.

25 **PMIX_GRPID** "pmix.egid" (uint32_t)
26 Effective group id.

27 -----

28 Host environments are not required to support any specific monitoring attributes.

▲-----▲

Optional Attributes

The following attributes may be implemented by a host environment.

- PMIX_MONITOR_ID** "pmix.monitor.id" (char*)
Provide a string identifier for this request.
- PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)
Identifier to be canceled (**NULL** means cancel all monitoring for this process).
- PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)
The application desires to control the response to a monitoring event.
- PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)
Register to have the PMIx server monitor the requestor for heartbeats.
- PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)
Time in seconds before declaring heartbeat missed.
- PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrops" (uint32_t)
Number of heartbeats that can be missed before generating the event.
- PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)
Register to monitor file for signs of life.
- PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)
Monitor size of given file is growing to determine if the application is running.
- PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)
Monitor time since last access of given file to determine if the application is running.
- PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
Monitor time since last modified of given file to determine if the application is running.
- PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
Time in seconds between checking the file.
- PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrops" (uint32_t)
Number of file checks that can be missed before generating the event.

Description

Request that a client be monitored for activity.

Advice to PMIx server hosts

If this module entry is provided and called by the PMIx server library, then the host environment must either provide the requested services or return **PMIX_ERR_NOT_SUPPORTED** to the provided *cbfunc*.

1 11.2.23 pmix_server_get_cred_fn_t

2 Summary

3 Request a credential from the host environment

4 Format

PMIx v3.0

```
▼----- C -----▼  
5 typedef pmix_status_t (*pmix_server_get_cred_fn_t) (  
6     const pmix_proc_t *proc,  
7     const pmix_info_t directives[],  
8     size_t ndirs,  
9     pmix_credential_cbfunc_t cbfunc,  
10    void *cbdata);  
▲----- C -----▲
```

11 IN proc

12 [pmix_proc_t](#) structure of requesting process (handle)

13 IN directives

14 Array of info structures (array of handles)

15 IN ndirs

16 Number of elements in the *info* array (integer)

17 IN cbfunc

18 Callback function to return the credential ([pmix_credential_cbfunc_t](#) function
19 reference)

20 IN cbdata

21 Data to be passed to the callback function (memory reference)

22 Returns [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant. In the event
23 the function returns an error, the *cbfunc* will not be called.

▼----- Required Attributes -----▼

24 If the PMIx library does not itself provide the requested credential, then it is required to pass any
25 attributes provided by the client to the host environment for processing. In addition, it must include
26 the following attributes in the passed *info* array:

27 **PMIX_USERID** "pmix.euid" ([uint32_t](#))

28 Effective user id.

29 **PMIX_GRPID** "pmix.egid" ([uint32_t](#))

30 Effective group id.
▲-----

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_CRED_TYPE "pmix.sec.ctype" (char*)

When passed in **PMIx_Get_credential**, a prioritized, comma-delimited list of desired credential types for use in environments where multiple authentication mechanisms may be available. When returned in a callback function, a string identifier of the credential type.

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Request a credential from the host environment

Advice to PMIx server hosts

If this module entry is provided and called by the PMIx server library, then the host environment must either provide the requested credential in the callback function or immediately return an error to the caller.

11.2.24 pmix_server_validate_cred_fn_t

Summary

Request validation of a credential

1
PMIx v3.0

Format

C

```
typedef pmix_status_t (*pmix_server_validate_cred_fn_t) (  
    const pmix_proc_t *proc,  
    const pmix_byte_object_t *cred,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_validation_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN `proc`
`pmix_proc_t` structure of requesting process (handle)

IN `cred`
Pointer to `pmix_byte_object_t` containing the credential (handle)

IN `directives`
Array of info structures (array of handles)

IN `ndirs`
Number of elements in the *info* array (integer)

IN `cbfunc`
Callback function to return the result (`pmix_validation_cbfunc_t` function reference)

IN `cbdata`
Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes


If the PMIx library does not itself validate the credential, then it is required to pass any attributes provided by the client to the host environment for processing. In addition, it must include the following attributes in the passed *info* array:

PMIX_USERID "pmix.euid" (`uint32_t`)
Effective user id.

PMIX_GRPID "pmix.egid" (`uint32_t`)
Effective group id.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19

Host environments are not required to support any specific attributes.



The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.



We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Request validation of a credential obtained from the host environment via a prior call to the `pmix_server_get_cred_fn_t` module entry.

11.2.25 pmix_server_iof_fn_t

Summary

Request the specified IO channels be forwarded from the given array of processes.

1
PMIx v3.0

Format

C

```
2 typedef pmix_status_t (*pmix_server_iof_fn_t) (  
3     const pmix_proc_t procs[], size_t nprocs,  
4     const pmix_info_t directives[], size_t ndirs,  
5     pmix_iof_channel_t channels,  
6     pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- 7 **IN** `procs`
8 Array `pmix_proc_t` identifiers whose IO is being requested (handle)
- 9 **IN** `nprocs`
10 Number of elements in `procs` (`size_t`)
- 11 **IN** `directives`
12 Array of `pmix_info_t` structures further defining the request (array of handles)
- 13 **IN** `ndirs`
14 Number of elements in the `info` array (integer)
- 15 **IN** `channels`
16 Bitmask identifying the channels to be forwarded (`pmix_iof_channel_t`)
- 17 **IN** `cbfunc`
18 Callback function `pmix_op_cbfunc_t` (function reference)
- 19 **IN** `cbdata`
20 Data to be passed to the callback function (memory reference)

21 Returns one of the following:

- 22 • `PMIX_SUCCESS` , indicating that the request is being processed by the host environment - result
23 will be returned in the provided `cbfunc`. Note that the library must not invoke the callback
24 function prior to returning from the API.
- 25 • `PMIX_OPERATION_SUCCEEDED` , indicating that the request was immediately processed and
26 returned `success` - the `cbfunc` will not be called
- 27 • a PMIx error constant indicating either an error in the input or that the request was immediately
28 processed and failed - the `cbfunc` will not be called

Required Attributes

29 The following attributes are required to be included in the passed `info` array:

- 30 **PMIX_USERID** "pmix.euid" (`uint32_t`)
31 Effective user id.
- 32 **PMIX_GRPID** "pmix.egid" (`uint32_t`)
33 Effective group id.

1
2
3
4
5
6
7
8
9
10
11

Host environments that provide this module entry point are required to support the following attributes:

- PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (uint32_t)
The requested size of the server cache in bytes for each specified channel. By default, the server is allowed (but not required) to drop all bytes received beyond the max size.
- PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (bool)
In an overflow situation, drop the oldest bytes to make room in the cache.
- PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (bool)
In an overflow situation, drop any new bytes received until room becomes available in the cache (default).

▲-----▲

▼----- Optional Attributes -----▼

12
13
14
15
16
17
18
19
20
21

- The following attributes may be supported by a host environment.
- PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (uint32_t)
Controls grouping of IO on the specified channel(s) to avoid being called every time a bit of IO arrives. The library will execute the callback whenever the specified number of bytes becomes available. Any remaining buffered data will be “flushed” upon call to deregister the respective channel.
 - PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (uint32_t)
Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering size, this prevents IO from being held indefinitely while waiting for another payload to arrive.



Description

Request the specified IO channels be forwarded from the given array of processes. An error shall be returned in the callback function if the requested service from any of the requested processes cannot be provided.

▼----- Advice to PMIx library implementers -----▼

The forwarding of stdin is a *push* process - processes cannot request that it be *pulled* from some other source. Requests including the **PMIX_FWD_STDIN_CHANNEL** channel will return a **PMIX_ERR_NOT_SUPPORTED** error.



1 11.2.26 pmix_server_stdin_fn_t

2 Summary

3 Pass standard input data to the host environment for transmission to specified recipients.

4 Format

PMIx v3.0

C

```
5 typedef pmix_status_t (*pmix_server_stdin_fn_t) (  
6     const pmix_proc_t *source,  
7     const pmix_proc_t targets[],  
8     size_t ntargets,  
9     const pmix_info_t directives[],  
10    size_t ndirs,  
11    const pmix_byte_object_t *bo,  
12    pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

13 IN **source**

14 [pmix_proc_t](#) structure of source process (handle)

15 IN **targets**

16 Array of [pmix_proc_t](#) target identifiers (handle)

17 IN **ntargets**

18 Number of elements in the *targets* array (integer)

19 IN **directives**

20 Array of info structures (array of handles)

21 IN **ndirs**

22 Number of elements in the *info* array (integer)

23 IN **bo**

24 Pointer to [pmix_byte_object_t](#) containing the payload (handle)

25 IN **cbfunc**

26 Callback function [pmix_op_cbfunc_t](#) (function reference)

27 IN **cbdata**

28 Data to be passed to the callback function (memory reference)

29 Returns one of the following:

- 30 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
31 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
32 function prior to returning from the API.
- 33 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
34 returned *success* - the *cbfunc* will not be called
- 35 • a PMIx error constant indicating either an error in the input or that the request was immediately
36 processed and failed - the *cbfunc* will not be called

Required Attributes

1 The following attributes are required to be included in the passed *info* array:

2 **PMIX_USERID** "pmix.euid" (uint32_t)

3 Effective user id.

4 **PMIX_GRPID** "pmix.egid" (uint32_t)

5 Effective group id.

Description

6 Passes stdin to the host environment for transmission to specified recipients. The host environment
7 is responsible for forwarding the data to all locations that host the specified *targets* and delivering
8 the payload to the PMIx server library connected to those clients.
9

Advice to PMIx server hosts

10 If this module entry is provided and called by the PMIx server library, then the host environment
11 must either provide the requested services or return **PMIX_ERR_NOT_SUPPORTED** to the
12 provided *cbfunc*.

CHAPTER 12

Scheduler-Specific Interfaces

The PMIx server library includes several interfaces specifically intended to support WLMs (also known as *schedulers*) by providing access to information of potential use to scheduling algorithms - e.g., information on communication costs between different points on the fabric. Due to their high cost in terms of execution, memory consumption, and interactions with other SMS components (e.g., a fabric manager), it is strongly advised that use be restricted to a single PMIx server in a system that is supporting the SMS component responsible for the scheduling of allocations (i.e., the system `scheduler`)

Accordingly, access to the functions described in this chapter requires that the PMIx server library be initialized with the `PMIX_SERVER_SCHEDULER` attribute.

12.1 Scheduler Support Datatypes

12.1.1 Fabric registration structure

The `pmix_fabric_t` structure is used by a WLM to interact with fabric-related PMIx interfaces.

PMIx v4.0

```
typedef struct pmix_fabric_s {
    char *name;
    size_t index;
    uint16_t **commcost;
    uint32_t nverts;
    void *module;
} pmix_fabric_t;
```

Note that in this structure:

- the *name* is an optional user-supplied string name identifying the fabric being referenced by this struct;
- a PMIx-provided index identifying this object;
- the *commcost* element is a square, two-dimensional array of `uint16_t` values representing the relative communication cost between the two (row,col) vertices. Note that PMIx makes no assumption as to the symmetry of the matrix - while the communication cost of many fabrics is independent of direction (and hence, the *commcost* matrix is symmetric), others may be direction sensitive;

- *nverts* indicates the number of rows and columns in the *commcost* array; and
- *module* points to an opaque object reserved for use by the PMIx server library.

The *name* field must be a **NULL**-terminated string composed of standard alphanumeric values supported by common utilities such as *strcmp*.

12.1.2 Scheduler Support Error Constants

PMIX_FABRIC_UPDATE_PENDING The PMIx server library has been alerted to a change in the fabric that requires updating of one or more registered `pmix_fabric_t` objects.

PMIX_FABRIC_UPDATED The PMIx server library has completed updating the entries of all affected `pmix_fabric_t` objects registered with the library. Access to the entries of those objects may now resume.

12.1.3 Scheduler Support Attributes

PMIX_SERVER_SCHEDULER "pmix.srv.sched" (bool)
Server requests access to WLM-supporting features

12.2 Scheduler Support Functions

The following APIs allow the scheduler that hosts the PMIx server library to request specific services from the PMIx library.

12.2.1 `PMIx_server_register_fabric`

Summary

Register for access to fabric-related information

Format

PMIx v4.0

```

pmix_status_t
PMIx_server_register_fabric(pmix_fabric_t *fabric,
                           const pmix_info_t directives[],
                           size_t ndirs)

```

IN fabric

address of a `pmix_fabric_t` (backed by storage). User may populate the "name" field at will - PMIx does not utilize this field (handle)

IN directives

an optional array of values indicating desired behaviors and/or fabric to be accessed. If **NULL**, then the highest priority available fabric will be used (array of handles)

IN ndirs

Number of elements in the *directives* array (integer)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_NETWORK_PLANE "pmix.net.plane" (char*)

ID string of a network plane (e.g., CIDR for Ethernet). When used as a modifier in a request for information, specifies the plane whose information is to be returned. When used directly in a request, returns a `pmix_data_array_t` of string identifiers for all network planes in the system.

Description

Register for access to fabric-related information, including communication cost matrix. This call must be made prior to requesting information from a fabric. The caller may request access to a particular `network plane` via the `PMIX_NETWORK_PLANE` attribute - otherwise, the default fabric will be returned.

If available, the *fabric* struct shall contain the address and size of the communication cost matrix associated with the specified network plane. For performance reasons, the PMIx server library does *not* provide thread protection for cost matrix access. Instead, users are required to register for `PMIX_FABRIC_UPDATE_PENDING` events indicating that an update to the cost matrix is pending. When received, users are required to terminate any actions involving access to the cost matrix before returning from the event.

Completion of the `PMIX_FABRIC_UPDATE_PENDING` event handler indicates to the PMIx server library that the fabric object's entries are available for updating. This may include releasing and re-allocating memory as the number of vertices may have changed (e.g., due to addition or removal of one or more NICs). When the update has been completed, the **PMIX! (PMIX!)** server library will generate a `PMIX_FABRIC_UPDATED` event indicating that it is safe to begin using the updated fabric object(s).

12.2.2 PMIx_server_deregister_fabric

Summary

Deregister a fabric object

Format

PMIx v4.0

```
pmix_status_t PMIx_server_deregister_fabric(pmix_fabric_t *fabric)
```

IN input

address of a `pmix_fabric_t` (handle)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

1 **Description**
2 Deregister a fabric object, providing an opportunity for the PMIx server library to cleanup any
3 information (e.g., cost matrix) associated with it.

4 **12.2.3 PMIx_server_get_vertex_info**

5 **Summary**
6 Given a communication cost matrix index for a specified fabric, return the corresponding vertex
7 info and the name of the node upon which it resides.

8 **Format**

PMIx v4.0

C

```
9 pmix_status_t  
10 PMIx_server_get_vertex_info(pmix_fabric_t *fabric,  
11                             uint32_t index, pmix_value_t *vertex,  
12                             char **nodename)
```

C

13 **IN fabric**
14 address of a `pmix_fabric_t` (handle)
15 **IN index**
16 communication cost matrix index (integer)
17 **IN vertex**
18 pointer to the `pmix_value_t` where the vertex info is to be returned (backed by storage)
19 (handle)
20 **OUT nodename**
21 pointer to the location where the string name of the host is to be returned. The caller is
22 responsible for releasing the string when done (handle)

23 Returns one of the following:

- 24 • **PMIX_SUCCESS**, indicating return of a valid value.
- 25 • **PMIX_ERR_BAD_PARAM**, indicating that the provided index is out of bounds
- 26 • a PMIx error constant indicating either an error in the input or that the request failed

27 **Description** 28 **12.2.4 PMIx_server_get_index**

29 **Summary**
30 Given vertex info, return the corresponding communication cost matrix index

1
PMIx v4.0

Format

C

```
2 pmix_status_t  
3 PMIx_server_get_index(pmix_fabric_t *fabric,  
4 pmix_value_t *vertex,  
5 uint32_t *index)
```

C

6 **IN fabric**
7 address of a `pmix_fabric_t` (handle)
8 **IN vertex**
9 pointer to the `pmix_value_t` containing the vertex info (handle)
10 **OUT index**
11 pointer to the location where the index is to be returned (memory reference (handle))

Description

Returns one of the following:

- `PMIX_SUCCESS`, indicating return of a valid value.
- a PMIx error constant indicating either an error in the input or that the request failed

Description

CHAPTER 13

Process Sets and Groups

PMIx supports two slightly related, but functionally different concepts known as *process sets* and *process groups*. This chapter these two concepts and describes how they are utilized, along with their corresponding APIs.

13.1 Process Sets

A PMIx *Process Set* is a user-provided label associated with a given set of application processes. Definition of a PMIx process set typically occurs at time of application execution - e.g., on a PRRTE command line:

```
$ prun -n 4 --pset ocean myoceanapp : -n 3 --pset ice myiceapp
```

In this example, the processes in the first application will be labeled with a **PMIX_PSET_NAME** attribute of *ocean* while those in the second application will be labeled with an *ice* value. During the execution, application processes could lookup the process set attribute for any other process using **PMIx_Get**. Alternatively, other executing applications could utilize the **PMIx_Query_info_nb** API to obtain the number of declared process sets in the system, a list of their names, and other information about them. In other words, the *process set* identifier provides a label by which an application can derive information about a process and its application - it does *not*, however, confer any operational function.

Thus, *process sets* differ from *process groups* in several key ways:

- *Process sets* have no implied relationship between their members - i.e., a process in a process set has no concept of a “pset rank” as it would in a process *group*
- Processes can only have one process *set* identifier, but can simultaneously belong to multiple process *groups*
- Process *set* identifiers are considered job-level information set at launch. No PMIx API is provided by which a user can change the process *set* value of a process on-the-fly. In contrast, PMIx process *groups* can only be defined dynamically by the application.

- Process *groups* can be used in calls to PMIx operations. Members of process *groups* that are involved in an operation are translated by their PMIx server into their *native* identifier prior to the operation being passed to the host environment. For example, an application can define a process group to consist of ranks 0 and 1 from the host-assigned namespace of 210456, identified by the group id of *foo*. If the application subsequently calls the **PMIx_Fence** API with a process identifier of {foo, PMIX_RANK_WILDCARD}, the PMIx server will replace that identifier with an array consisting of {210456, 0} and {210456, 1} - the host-assigned identifiers of the participating processes - prior to passing the request up to the host environment
- Process *groups* can request that the host environment assign a unique **size_t** PGCID to the group at time of group construction. An MPI library may, for example, use the PGCID as the MPI communicator identifier for the group.

The two concepts do, however, overlap in one specific area. Process *groups* are included in the process *set* information returned by calls to **PMIx_Query_info_nb**. Thus, a *process group* can effectively be considered an extended version of a *process set* that adds dynamic definition and operational context to the *process set* concept.

Advice to PMIx library implementers

PMIx implementations are required to include all active *group* identifiers in the returned list of process *set* names provided in response to the appropriate **PMIx_Query_info_nb** call.

13.2 Process Groups

PMIx *Groups* are defined as a collection of processes desiring a common, unique identifier for purposes such as passing events or participating in PMIx fence operations. As with processes that assemble via **PMIx_Connect**, each member of the group is provided with both the job-level information of any other namespace represented in the group, and the contact information for all group members. However, *groups* differ from **PMIx_Connect** assemblages in the following key areas:

- Relation to the host environment
 - Calls to **PMIx_Connect** are relayed to the host environment. This means that the host RM should treat the failure of any process in the specified assemblage as a reportable event and take appropriate action. However, the environment is not required to define a new identifier for the connected assemblage or any of its member processes, nor does it define a new rank for each process within that assemblage. In addition, the PMIx server does not provide any tracking support for the assemblage. Thus, the caller is responsible for addressing members of the connected assemblage using their RM-provided identifiers.

- 1 – Calls to PMIx Group APIs are first processed within the local PMIx server. When constructed,
2 the server creates a tracker that associates the specified processes with the user-provided group
3 identifier, and assigns a new *group rank* based on their relative position in the array of
4 processes provided in the call to `PMIx_Group_construct` . Members of the group can
5 subsequently utilize the group identifier in PMIx function calls to address the group’s
6 members, using either `PMIX_RANK_WILDCARD` to refer to all of them or the group-level
7 rank of specific members. The PMIx server will translate the specified processes into their
8 RM-assigned identifiers prior to passing the request up to its host. Thus, the host environment
9 has no visibility into the group’s existence or membership.

Advice to users

10 User-provided group identifiers must be distinct from anything provided by the RM so as to
11 avoid collisions between group identifiers and RM-assigned namespaces. This can usually be
12 accomplished through the use of an application-specific prefix – e.g., “myapp-foo”

• Construction procedure

- 13 – `PMIx_Connect` calls require that every process call the API before completing – i.e., it is
14 modeled upon the bulk synchronous traditional MPI connect/accept methodology. Thus, a
15 given application thread can only be involved in one connect/accept operation at a time, and is
16 blocked in that operation until all specified processes participate. In addition, there is no
17 provision for replacing processes in the assemblage due to failure to participate, nor a
18 mechanism by which a process might decline participation.
19 – PMIx Groups are designed to be more flexible in their construction procedure by relaxing
20 these constraints. While a standard blocking form of constructing groups is provided, the event
21 notification system is utilized to provide a designated *group leader* with the ability to replace
22 participants that fail to participate within a given timeout period. This provides a mechanism
23 by which the application can, if desired, replace members on-the-fly or allow the group to
24 proceed with partial membership. In such cases, the final group membership is returned to all
25 participants upon completion of the operation.
26 – PMIx Groups are designed to be more flexible in their construction procedure by relaxing
27 these constraints. While a standard blocking form of constructing groups is provided, the event
28 notification system is utilized to provide a designated *group leader* with the ability to replace
29 participants that fail to participate within a given timeout period. This provides a mechanism
30 by which the application can, if desired, replace members on-the-fly or allow the group to
31 proceed with partial membership. In such cases, the final group membership is returned to all
32 participants upon completion of the operation.

27 Additionally, PMIx supports dynamic definition of group membership based on an invite/join
28 model. A process can asynchronously initiate construction of a group of any processes via the
29 `PMIx_Group_invite` function call. Invitations are delivered via a PMIx event (using the
30 `PMIX_GROUP_INVITED` event) to the invited processes which can then either accept or
31 decline the invitation using the `PMIx_Group_join` API. The initiating process tracks
32 responses by registering for the events generated by the call to `PMIx_Group_join` ,
33 timeouts, or process terminations, optionally replacing processes that decline the invitation,
34 fail to respond in time, or terminate without responding. Upon completion of the operation,
35 the final list of participants is communicated to each member of the new group.

• Destruct procedure

- Processes that assemble via **PMIx_Connect** must all depart the assemblage together – i.e., no member can depart the assemblage while leaving the remaining members in it. Even the non-blocking form of **PMIx_Disconnect** retains this requirement in that members remain a part of the assemblage until all members have called **PMIx_Disconnect_nb**
- Members of a PMIx Group may depart the group at any time via the **PMIx_Group_leave** API. Other members are notified of the departure via the **PMIX_GROUP_LEFT** event to distinguish such events from those reporting process termination. This leaves the remaining members free to continue group operations. The **PMIx_Group_destruct** operation offers a collective method akin to **PMIx_Disconnect** for deconstructing the entire group.

Note that applications supporting dynamic group behaviors such as asynchronous departure take responsibility for ensuring global consistency in the group definition prior to executing group collective operations - i.e., it is the application's responsibility to either ensure that knowledge of the current group membership is globally consistent across the participants, or to register for appropriate events to deal with the lack of consistency during the operation.

In other words, members of PMIx Groups are *loosely coupled* as opposed to *tightly connected* when constructed via **PMIx_Connect** . The relevant APIs are explained below.

Advice to users

The reliance on PMIx events in the PMIx Group concept dictates that processes utilizing these APIs must register for the corresponding events. Failure to do so will likely lead to operational failures. Users are recommended to utilize the **PMIX_TIMEOUT** directive (or retain an internal timer) on calls to PMIx Group APIs (especially the blocking form of those functions) as processes that have not registered for required events will never respond.

13.2.1 PMIx_Group_construct

Summary

Construct a PMIx process group

Format

C

```

2 pmix_status_t
3 PMIx_Group_construct(const char grp[],
4                     const pmix_proc_t procs[], size_t nprocs,
5                     const pmix_info_t directives[], size_t ndirs,
6                     pmix_info_t **results, size_t *nresults)

```

C

- 7 **IN grp**
8 NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the
9 group identifier (string)
- 10 **IN procs**
11 Array of **pmix_proc_t** structures containing the PMIx identifiers of the member processes
12 (array of handles)
- 13 **IN nprocs**
14 Number of elements in the *procs* array (**size_t**)
- 15 **IN directives**
16 Array of **pmix_info_t** structures (array of handles)
- 17 **IN ndirs**
18 Number of elements in the *directives* array (**size_t**)
- 19 **INOUT results**
20 Pointer to a location where the array of **pmix_info_t** describing the results of the
21 operation is to be returned (pointer to handle)
- 22 **INOUT nresults**
23 Pointer to a **size_t** location where the number of elements in *results* is to be returned
24 (memory reference)

25 Returns one of the following:

- 26 • **PMIX_SUCCESS** , indicating that the request has been successfully completed
- 27 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library and/or the host RM does not support this
28 operation
- 29 • a PMIx error constant indicating either an error in the input or that the request failed to be
30 completed

Required Attributes

31 The following attributes are *required* to be supported by all PMIx libraries that support this
32 operation:

- 33 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)
34 This process is the leader of the group
- 35 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

1 Participation is optional - do not return an error if any of the specified processes terminate
2 without having joined. The default is false

3 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)

4 Group operation only involves local processes. PMIx implementations are *required* to
5 automatically scan an array of group members for local vs remote processes - if only local
6 processes are detected, the implementation need not execute a global collective for the
7 operation unless a context ID has been requested from the host environment. This can result
8 in significant time savings. This attribute can be used to optimize the operation by indicating
9 whether or not only local processes are represented, thus allowing the implementation to
10 bypass the scan. The default is false

11 Host environments that support this operation are *required* to provide the following attributes:

12 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

13 Requests that the RM assign a new context identifier to the newly created group. The
14 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
15 specified in the request. Thus, the value serves as a means of identifying the group within
16 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION** .

17 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

18 Notify remaining members when another member terminates without first leaving the group.
19 The default is false



Optional Attributes



20 The following attributes are optional for host environments that support this operation:

21 **PMIX_TIMEOUT** "pmix.timeout" (int)

22 Time in seconds before the specified operation should time out (0 indicating infinite) in
23 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
24 the target process from ever exposing its data.



Advice to PMIx library implementers

25 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
26 environment due to race condition considerations between completion of the operation versus
27 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
28 directly in the PMIx server library must take care to resolve the race condition and should avoid
29 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
30 created.



Description

Construct a new group composed of the specified processes and identified with the provided group identifier. The group identifier is a user-defined, **NULL**-terminated character array of length less than or equal to **PMIX_MAX_NSLEN** . Only characters accepted by standard string comparison functions (e.g., *strncmp*) are supported. Processes may engage in multiple simultaneous group construct operations so long as each is provided with a unique group ID. The *directives* array can be used to pass user-level directives regarding timeout constraints and other options available from the PMIx server.

If the **PMIX_GROUP_NOTIFY_TERMINATION** attribute is provided and has a value of **true**, then either the construct leader (if **PMIX_GROUP_LEADER** is provided) or all participants who register for the **PMIX_GROUP_MEMBER_FAILED** event will receive events whenever a process fails or terminates prior to calling **PMIx_Group_construct** – i.e. if a *group leader* is declared, *only* that process will receive the event. In the absence of a declared leader, *all* specified group members will receive the event.

The event will contain the identifier of the process that failed to join plus any other information that the host RM provided. This provides an opportunity for the leader or the collective members to react to the event – e.g., to decide to proceed with a smaller group or to abort the operation. The decision is communicated to the PMIx library in the results array at the end of the event handler. This allows PMIx to properly adjust accounting for procedure completion. When construct is complete, the participating PMIx servers will be alerted to any change in participants and each group member will receive an updated group membership (marked with the **PMIX_GROUP_MEMBERSHIP** attribute) as part of the *results* array returned by this API.

Failure of the declared leader at any time will cause a **PMIX_GROUP_LEADER_FAILED** event to be delivered to all participants so they can optionally declare a new leader. A new leader is identified by providing the **PMIX_GROUP_LEADER** attribute in the results array in the return of the event handler. Only one process is allowed to return that attribute, thereby declaring itself as the new leader. Results of the leader selection will be communicated to all participants via a **PMIX_GROUP_LEADER_SELECTED** event identifying the new leader. If no leader was selected, then the **pmix_info_t** provided to that event handler will include that information so the participants can take appropriate action.

Any participant that returns **PMIX_GROUP_CONSTRUCT_ABORT** from either the **PMIX_GROUP_MEMBER_FAILED** or the **PMIX_GROUP_LEADER_FAILED** event handler will cause the construct process to abort, returning from the call with a **PMIX_GROUP_CONSTRUCT_ABORT** status.

If the **PMIX_GROUP_NOTIFY_TERMINATION** attribute is not provided or has a value of **false**, then the **PMIx_Group_construct** operation will simply return an error whenever a proposed group member fails or terminates prior to calling **PMIx_Group_construct** .

Providing the **PMIX_GROUP_OPTIONAL** attribute with a value of **true** directs the PMIx library to consider participation by any specified group member as non-required - thus, the operation will return **PMIX_SUCCESS** if all members participate, or **PMIX_ERR_PARTIAL_SUCCESS** if

1 some members fail to participate. The *results* array will contain the final group membership in the
2 latter case. Note that this use-case can cause the operation to hang if the `PMIX_TIMEOUT`
3 attribute is not specified and one or more group members fail to call `PMIX_Group_construct`
4 while continuing to execute. Also, note that no leader or member failed events will be generated
5 during the operation.

6 Processes in a group under construction are not allowed to leave the group until group construction
7 is complete. Upon completion of the construct procedure, each group member will have access to
8 the job-level information of all namespaces represented in the group plus any information posted
9 via `PMIX_Put` (subject to the usual scoping directives) for every group member.

▼ ————— Advice to PMIx library implementers ————— ▼

10 At the conclusion of the construct operation, the PMIx library is *required* to ensure that job-related
11 information from each participating namespace plus any information posted by group members via
12 `PMIX_Put` (subject to scoping directives) is available to each member via calls to `PMIX_Get` .



▼ ————— Advice to PMIx server hosts ————— ▼

13 The collective nature of this API generally results in use of a fence-like operation by the backend
14 host environment. Host environments that utilize the array of process participants as a *signature* for
15 such operations may experience potential conflicts should both a `PMIX_Group_construct`
16 and a `PMIX_Fence` operation involving the same participants be simultaneously executed. As
17 PMIx allows for such use-cases, it is therefore the responsibility of the host environment to resolve
18 any potential conflicts.



19 13.2.2 `PMIX_Group_construct_nb`

20 Summary

21 Non-blocking form of `PMIX_Group_construct`

1
PMIx v4.0

Format

C

```
2 pmix_status_t
3 PMIx_Group_construct_nb(const char grp[],
4                         const pmix_proc_t procs[], size_t nprocs,
5                         const pmix_info_t directives[], size_t ndirs,
6                         pmix_info_cbfunc_t cbfunc, void *cbdata)
```

C

- 7 **IN** **grp**
8 NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the
9 group identifier (string)
- 10 **IN** **procs**
11 Array of **pmix_proc_t** structures containing the PMIx identifiers of the member processes
12 (array of handles)
- 13 **IN** **nprocs**
14 Number of elements in the *procs* array (**size_t**)
- 15 **IN** **directives**
16 Array of **pmix_info_t** structures (array of handles)
- 17 **IN** **ndirs**
18 Number of elements in the *directives* array (**size_t**)
- 19 **IN** **cbfunc**
20 Callback function **pmix_info_cbfunc_t** (function reference)
- 21 **IN** **cbdata**
22 Data to be passed to the callback function (memory reference)

23 Returns one of the following:

- 24 • **PMIX_SUCCESS** indicating that the request has been accepted for processing and the provided
25 callback function will be executed upon completion of the operation. Note that the library *must*
26 *not* invoke the callback function prior to returning from the API.
- 27 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
28 returned *success* - the *cbfunc* will *not* be called
- 29 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library does not support this operation - the *cbfunc*
30 will *not* be called
- 31 • a non-zero PMIx error constant indicating a reason for the request to have been rejected - the
32 *cbfunc* will *not* be called

33 If executed, the status returned in the provided callback function will be one of the following
34 constants:

- 35 • **PMIX_SUCCESS** The operation succeeded and all specified members participated.

- 1 • **PMIX_ERR_PARTIAL_SUCCESS** The operation succeeded but not all specified members
2 participated - the final group membership is included in the callback function
- 3 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
4 does not.
- 5 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

6 PMIx libraries that choose not to support this operation *must* return
7 **PMIX_ERR_NOT_SUPPORTED** when the function is called.

8 The following attributes are *required* to be supported by all PMIx libraries that support this
9 operation:

10 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)

11 This process is the leader of the group

12 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

13 Participation is optional - do not return an error if any of the specified processes terminate
14 without having joined. The default is false

15 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)

16 Group operation only involves local processes. PMIx implementations are *required* to
17 automatically scan an array of group members for local vs remote processes - if only local
18 processes are detected, the implementation need not execute a global collective for the
19 operation unless a context ID has been requested from the host environment. This can result
20 in significant time savings. This attribute can be used to optimize the operation by indicating
21 whether or not only local processes are represented, thus allowing the implementation to
22 bypass the scan. The default is false

23 Host environments that support this operation are *required* to provide the following attributes:

24 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

25 Requests that the RM assign a new context identifier to the newly created group. The
26 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
27 specified in the request. Thus, the value serves as a means of identifying the group within
28 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

29 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

30 Notify remaining members when another member terminates without first leaving the group.
31 The default is false

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking version of the **PMIx_Group_construct** operation. The callback function will be called once all group members have called either **PMIx_Group_construct** or **PMIx_Group_construct_nb**.

13.2.3 PMIx_Group_destruct

Summary

Destruct a PMIx process group

1
PMIx v4.0

Format

C

```

2 pmix_status_t
3 PMIx_Group_destruct(const char grp[],
4                     const pmix_info_t directives[], size_t ndirs)

```

C

- 5 **IN** `grp`
- 6 NULL-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the
- 7 identifier of the group to be destructed (string)
- 8 **IN** `directives`
- 9 Array of `pmix_info_t` structures (array of handles)
- 10 **IN** `ndirs`
- 11 Number of elements in the `directives` array (`size_t`)

12 Returns one of the following:

- 13 • `PMIX_SUCCESS`, indicating that the request has been successfully completed
- 14 • `PMIX_ERR_NOT_SUPPORTED` The PMIx library and/or the host RM does not support this
- 15 operation
- 16 • a PMIx error constant indicating either an error in the input or that the request failed to be
- 17 completed

Required Attributes

18 For implementations and host environments that support the operation, there are no identified
19 required attributes for this API.

Optional Attributes

20 The following attributes are optional for host environments that support this operation:

- 21 `PMIX_TIMEOUT` "`pmix.timeout`" (`int`)
- 22 Time in seconds before the specified operation should time out (0 indicating infinite) in
- 23 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
- 24 the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not created.

Description

Destruct a group identified by the provided group identifier. Processes may engage in multiple simultaneous group destruct operations so long as each involves a unique group ID. The *directives* array can be used to pass user-level directives regarding timeout constraints and other options available from the PMIx server.

The destruct API will return an error if any group process fails or terminates prior to calling `PMIx_Group_destruct` or its non-blocking version unless the `PMIX_GROUP_NOTIFY_TERMINATION` attribute was provided (with a value of `false`) at time of group construction. If notification was requested, then the `PMIX_GROUP_MEMBER_FAILED` event will be delivered for each process that fails to call destruct and the destruct tracker updated to account for the lack of participation. The `PMIx_Group_destruct` operation will subsequently return `PMIX_SUCCESS` when the remaining processes have all called destruct – i.e., the event will serve in place of return of an error.

Advice to PMIx server hosts

The collective nature of this API generally results in use of a fence-like operation by the backend host environment. Host environments that utilize the array of process participants as a *signature* for such operations may experience potential conflicts should both a `PMIx_Group_destruct` and a `PMIx_Fence` operation involving the same participants be simultaneously executed. As PMIx allows for such use-cases, it is therefore the responsibility of the host environment to resolve any potential conflicts.

13.2.4 `PMIx_Group_destruct_nb`

Summary

Non-blocking form of `PMIx_Group_destruct`

1
PMIx v4.0

Format

C

```
2 pmix_status_t
3 PMIx_Group_destruct_nb(const char grp[],
4                       const pmix_info_t directives[], size_t ndirs,
5                       pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

- 6 **IN grp**
7 NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the
8 identifier of the group to be destructed (string)
- 9 **IN directives**
10 Array of **pmix_info_t** structures (array of handles)
- 11 **IN ndirs**
12 Number of elements in the *directives* array (**size_t**)
- 13 **IN cbfunc**
14 Callback function **pmix_op_cbfunc_t** (function reference)
- 15 **IN cbdata**
16 Data to be passed to the callback function (memory reference)

17 Returns one of the following:

- 18 • **PMIX_SUCCESS**, indicating that the request is being processed - result will be returned in the
19 provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning
20 from the API.
- 21 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
22 returned *success* - the *cbfunc* will *not* be called
- 23 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library does not support this operation - the *cbfunc*
24 will *not* be called
- 25 • a PMIx error constant indicating either an error in the input or that the request was immediately
26 processed and failed - the *cbfunc* will *not* be called

27 If executed, the status returned in the provided callback function will be one of the following
28 constants:

- 29 • **PMIX_SUCCESS** The operation was successfully completed
- 30 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
31 does not.
- 32 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

PMIx libraries that choose not to support this operation *must* return `PMIX_ERR_NOT_SUPPORTED` when the function is called. For implementations and host environments that support the operation, there are no identified required attributes for this API.

Optional Attributes

The following attributes are optional for host environments that support this operation:

`PMIX_TIMEOUT` "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking version of the `PMIx_Group_destruct` operation. The callback function will be called once all members of the group have executed either `PMIx_Group_destruct` or `PMIx_Group_destruct_nb`.

13.2.5 PMIx_Group_invite

Summary

Asynchronously construct a PMIx process group

1
PMIx v4.0

Format

C

```

2 pmix_status_t
3 PMIx_Group_invite(const char grp[],
4                   const pmix_proc_t procs[], size_t nprocs,
5                   const pmix_info_t directives[], size_t ndirs,
6                   pmix_info_t **results, size_t *nresult)

```

C

- 7 **IN grp**
8 NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the
9 group identifier (string)
- 10 **IN procs**
11 Array of **pmix_proc_t** structures containing the PMIx identifiers of the processes to be
12 invited (array of handles)
- 13 **IN nprocs**
14 Number of elements in the *procs* array (**size_t**)
- 15 **IN directives**
16 Array of **pmix_info_t** structures (array of handles)
- 17 **IN ndirs**
18 Number of elements in the *directives* array (**size_t**)
- 19 **INOUT results**
20 Pointer to a location where the array of **pmix_info_t** describing the results of the
21 operation is to be returned (pointer to handle)
- 22 **INOUT nresults**
23 Pointer to a **size_t** location where the number of elements in *results* is to be returned
24 (memory reference)

25 Returns one of the following:

- 26 • **PMIX_SUCCESS** , indicating that the request has been successfully completed
- 27 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library and/or the host RM does not support this
28 operation
- 29 • a PMIx error constant indicating either an error in the input or that the request failed to be
30 completed

Required Attributes

31 The following attributes are *required* to be supported by all PMIx libraries that support this
32 operation:

- 33 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (**bool**)
34 Participation is optional - do not return an error if any of the specified processes terminate
35 without having joined. The default is false

1 Host environments that support this operation are *required* to provide the following attributes:

2 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

3 Requests that the RM assign a new context identifier to the newly created group. The
4 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
5 specified in the request. Thus, the value serves as a means of identifying the group within
6 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

7 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

8 Notify remaining members when another member terminates without first leaving the group.
9 The default is false

▲-----▲
▼-----▼ **Optional Attributes** -----▼

10 The following attributes are optional for host environments that support this operation:

11 **PMIX_TIMEOUT** "pmix.timeout" (int)

12 Time in seconds before the specified operation should time out (0 indicating infinite) in
13 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
14 the target process from ever exposing its data.

▲-----▲
▼-----▼ **Advice to PMIx library implementers** -----▼

15 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
16 environment due to race condition considerations between completion of the operation versus
17 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
18 directly in the PMIx server library must take care to resolve the race condition and should avoid
19 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
20 created.

Description

Explicitly invite the specified processes to join a group. The process making the `PMIx_Group_invite` call is automatically declared to be the *group leader*. Each invited process will be notified of the invitation via the `PMIX_GROUP_INVITED` event - the processes being invited must therefore register for the `PMIX_GROUP_INVITED` event in order to be notified of the invitation. Note that the PMIx event notification system caches events - thus, no ordering of invite versus event registration is required.

The invitation event will include the identity of the inviting process plus the name of the group. When ready to respond, each invited process provides a response using either the blocking or non-blocking form of `PMIx_Group_join`. This will notify the inviting process that the invitation was either accepted (via the `PMIX_GROUP_INVITE_ACCEPTED` event) or declined (via the `PMIX_GROUP_INVITE_DECLINED` event). The `PMIX_GROUP_INVITE_ACCEPTED` event is captured by the PMIx client library of the inviting process - i.e., the application itself does not need to register for this event. The library will track the number of accepting processes and alert the inviting process (by returning from the blocking form of `PMIx_Group_invite` or calling the callback function of the non-blocking form) when group construction completes.

The inviting process should, however, register for the `PMIX_GROUP_INVITE_DECLINED` if the application allows invited processes to decline the invitation. This provides an opportunity for the application to either invite a replacement, declare “abort”, or choose to remove the declining process from the final group. The inviting process should also register to receive `PMIX_GROUP_INVITE_FAILED` events whenever a process fails or terminates prior to responding to the invitation. Actions taken by the inviting process in response to these events must be communicated at the end of the event handler by returning the corresponding result so that the PMIx library can adjust accordingly.

Upon completion of the operation, all members of the new group will receive access to the job-level information of each other’s namespaces plus any information posted via `PMIx_Put` by the other members.

The inviting process is automatically considered the leader of the asynchronous group construction procedure and will receive all failure or termination events for invited members prior to completion. The inviting process is required to provide a `PMIX_GROUP_CONSTRUCT_COMPLETE` event once the group has been fully assembled - this event is used by the PMIx library as a trigger to release participants from their call to `PMIx_Group_join` and provides information (e.g., the final group membership) to be returned in the *results* array.

Advice to users

Applications are not allowed to use the group in any operations until group construction is complete. This is required in order to ensure consistent knowledge of group membership across all participants.

Failure of the inviting process at any time will cause a `PMIX_GROUP_LEADER_FAILED` event to be delivered to all participants so they can optionally declare a new leader. A new leader is identified by providing the `PMIX_GROUP_LEADER` attribute in the results array in the return of the event handler. Only one process is allowed to return that attribute, declaring itself as the new leader. Results of the leader selection will be communicated to all participants via a `PMIX_GROUP_LEADER_SELECTED` event identifying the new leader. If no leader was selected, then the status code provided in the event handler will provide an error value so the participants can take appropriate action.

13.2.6 `PMIx_Group_invite_nb`

Summary

Non-blocking form of `PMIx_Group_invite`

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_invite_nb(const char grp[],
                    const pmix_proc_t procs[], size_t nprocs,
                    const pmix_info_t directives[], size_t ndirs,
                    pmix_info_cbfunc_t cbfunc, void *cbdata)
```

C

- IN** `grp`
NULL-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group identifier (string)
- IN** `procs`
Array of `pmix_proc_t` structures containing the PMIx identifiers of the processes to be invited (array of handles)
- IN** `nprocs`
Number of elements in the `procs` array (`size_t`)
- IN** `directives`
Array of `pmix_info_t` structures (array of handles)
- IN** `ndirs`
Number of elements in the `directives` array (`size_t`)
- IN** `cbfunc`
Callback function `pmix_info_cbfunc_t` (function reference)
- IN** `cbdata`
Data to be passed to the callback function (memory reference)

Returns one of the following:

- 1 • **PMIX_SUCCESS** , indicating that the request is being processed - result will be returned in the
2 provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning
3 from the API.
- 4 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
5 returned *success* - the *cbfunc* will *not* be called
- 6 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library does not support this operation - the *cbfunc*
7 will *not* be called
- 8 • a PMIx error constant indicating either an error in the input or that the request was immediately
9 processed and failed - the *cbfunc* will *not* be called

10 If executed, the status returned in the provided callback function will be one of the following
11 constants:

- 12 • **PMIX_SUCCESS** The operation succeeded and all specified members participated.
- 13 • **PMIX_ERR_PARTIAL_SUCCESS** The operation succeeded but not all specified members
14 participated - the final group membership is included in the callback function
- 15 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
16 does not.
- 17 • a non-zero PMIx error constant indicating a reason for the request's failure

Required Attributes

18 The following attributes are *required* to be supported by all PMIx libraries that support this
19 operation:

20 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

21 Participation is optional - do not return an error if any of the specified processes terminate
22 without having joined. The default is false

23 Host environments that support this operation are *required* to provide the following attributes:

24 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

25 Requests that the RM assign a new context identifier to the newly created group. The
26 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
27 specified in the request. Thus, the value serves as a means of identifying the group within
28 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION** .

29 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

30 Notify remaining members when another member terminates without first leaving the group.
31 The default is false

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (0 indicating infinite) in error. The timeout parameter can help avoid “hangs” due to programming errors that prevent the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT** directly in the PMIx server library must take care to resolve the race condition and should avoid passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking version of the **PMIx_Group_invite** operation. The callback function will be called once all invited members of the group (or their substitutes) have executed either **PMIx_Group_join** or **PMIx_Group_join_nb**.

13.2.7 PMIx_Group_join

Summary

Accept an invitation to join a PMIx process group

1
PMIx v4.0

Format

C

```

2 pmix_status_t
3 PMIx_Group_join(const char grp[],
4                 const pmix_proc_t *leader,
5                 pmix_group_opt_t opt,
6                 const pmix_info_t directives[], size_t ndirs,
7                 pmix_info_t **results, size_t *nresult)

```

C

- 8 **IN grp**
NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group identifier (string)
- 10 **IN leader**
Process that generated the invitation (handle)
- 12 **IN opt**
Accept or decline flag (**pmix_group_opt_t**)
- 14 **IN directives**
Array of **pmix_info_t** structures (array of handles)
- 16 **IN ndirs**
Number of elements in the *directives* array (**size_t**)
- 18 **INOUT results**
Pointer to a location where the array of **pmix_info_t** describing the results of the operation is to be returned (pointer to handle)
- 20 **INOUT nresults**
Pointer to a **size_t** location where the number of elements in *results* is to be returned (memory reference)

25 Returns one of the following:

- 26 • **PMIX_SUCCESS**, indicating that the request has been successfully completed
- 27 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library and/or the host RM does not support this operation
- 28 • a PMIx error constant indicating either an error in the input or that the request failed to be completed

Required Attributes

31 There are no identified required attributes for implementers.

Optional Attributes

1 The following attributes are optional for host environments that support this operation:

2 **PMIX_TIMEOUT** "pmix.timeout" (int)

3 Time in seconds before the specified operation should time out (0 indicating infinite) in
4 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
5 the target process from ever exposing its data.

Advice to PMIx library implementers

6 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
7 environment due to race condition considerations between completion of the operation versus
8 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
9 directly in the PMIx server library must take care to resolve the race condition and should avoid
10 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
11 created.

Description

12 Respond to an invitation to join a group that is being asynchronously constructed. The process must
13 have registered for the **PMIX_GROUP_INVITED** event in order to be notified of the invitation.
14 When called, the event information will include the **pmix_proc_t** identifier of the process that
15 generated the invitation along with the identifier of the group being constructed. When ready to
16 respond, the process provides a response using either form of **PMIx_Group_join** .
17

Advice to users

18 Since the process is alerted to the invitation in a PMIx event handler, the process *must not* use the
19 blocking form of this call unless it first “thread shifts” out of the handler and into its own thread
20 context. Likewise, while it is safe to call the non-blocking form of the API from the event handler,
21 the process *must not* block in the handler while waiting for the callback function to be called.

1 Calling this function causes the inviting process (aka the *group leader*) to be notified that the
2 process has either accepted or declined the request. The blocking form of the API will return once
3 the group has been completely constructed or the group's construction has failed (as described
4 below) – likewise, the callback function of the non-blocking form will be executed upon the same
5 conditions.

6 Failure of the leader during the call to `PMIx_Group_join` will cause a
7 `PMIX_GROUP_LEADER_FAILED` event to be delivered to all invited participants so they can
8 optionally declare a new leader. A new leader is identified by providing the
9 `PMIX_GROUP_LEADER` attribute in the results array in the return of the event handler. Only one
10 process is allowed to return that attribute, declaring itself as the new leader. Results of the leader
11 selection will be communicated to all participants via a `PMIX_GROUP_LEADER_SELECTED`
12 event identifying the new leader. If no leader was selected, then the status code provided in the
13 event handler will provide an error value so the participants can take appropriate action.

14 Any participant that returns `PMIX_GROUP_CONSTRUCT_ABORT` from the leader failed event
15 handler will cause all participants to receive an event notifying them of that status. Similarly, the
16 leader may elect to abort the procedure by either returning `PMIX_GROUP_CONSTRUCT_ABORT`
17 from the handler assigned to the `PMIX_GROUP_INVITE_ACCEPTED` or
18 `PMIX_GROUP_INVITE_DECLINED` codes, or by generating an event for the abort code. Abort
19 events will be sent to all invited participants.

20 13.2.8 `PMIx_Group_join_nb`

21 Summary

22 Non-blocking form of `PMIx_Group_join`

23 Format

PMIx v4.0

```
24 pmix_status_t  
25 PMIx_Group_join_nb(const char grp[],  
26                   const pmix_proc_t *leader,  
27                   pmix_group_opt_t opt,  
28                   const pmix_info_t directives[], size_t ndirs,  
29                   pmix_info_cbfunc_t cbfunc, void *cbdata)
```

30 IN `grp`

31 `NULL`-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the
32 group identifier (string)

33 IN `leader`

34 Process that generated the invitation (handle)

1 **IN** **opt**
2 Accept or decline flag (`pmix_group_opt_t`)
3 **IN** **directives**
4 Array of `pmix_info_t` structures (array of handles)
5 **IN** **ndirs**
6 Number of elements in the *directives* array (`size_t`)
7 **IN** **cbfunc**
8 Callback function `pmix_info_cbfunc_t` (function reference)
9 **IN** **cbdata**
10 Data to be passed to the callback function (memory reference)

11 Returns one of the following:

- 12 • **PMIX_SUCCESS** , indicating that the request is being processed - result will be returned in the
13 provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning
14 from the API.
- 15 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
16 returned *success* - the *cbfunc* will *not* be called
- 17 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library does not support this operation - the *cbfunc*
18 will *not* be called
- 19 • a PMIx error constant indicating either an error in the input or that the request was immediately
20 processed and failed - the *cbfunc* will *not* be called

21 If executed, the status returned in the provided callback function will be one of the following
22 constants:

- 23 • **PMIX_SUCCESS** The operation succeeded and group membership is in the callback function
24 parameters
- 25 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
26 does not.
- 27 • a non-zero PMIx error constant indicating a reason for the request's failure

▼----- Required Attributes -----▼

28 There are no identified required attributes for implementers.

▲-----

▼----- Optional Attributes -----▼

29 The following attributes are optional for host environments that support this operation:

30 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)
31 Time in seconds before the specified operation should time out (0 indicating infinite) in
32 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
33 the target process from ever exposing its data.

Advice to PMIx library implementers

We recommend that implementation of the `PMIX_TIMEOUT` attribute be left to the host environment due to race condition considerations between completion of the operation versus internal timeout in the PMIx server library. Implementers that choose to support `PMIX_TIMEOUT` directly in the PMIx server library must take care to resolve the race condition and should avoid passing `PMIX_TIMEOUT` to the host environment so that multiple competing timeouts are not created.

Description

Non-blocking version of the `PMIx_Group_join` operation. The callback function will be called once all invited members of the group (or their substitutes) have executed either `PMIx_Group_join` or `PMIx_Group_join_nb`.

13.2.9 PMIx_Group_leave

Summary

Leave a PMIx process group

Format

PMIx v4.0

```
pmix_status_t
PMIx_Group_leave(const char grp[],
                 const pmix_info_t directives[], size_t ndirs)
```

IN grp

NULL-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group identifier (string)

IN directives

Array of `pmix_info_t` structures (array of handles)

IN ndirs

Number of elements in the *directives* array (`size_t`)

Returns one of the following:

- `PMIX_SUCCESS`, indicating that the request has been communicated to the local PMIx server
- `PMIX_ERR_NOT_SUPPORTED` The PMIx library and/or the host RM does not support this operation
- a PMIx error constant indicating either an error in the input or that the request is unsupported

Required Attributes

There are no identified required attributes for implementers.

Description

Leave a PMIx Group. Calls to **PMIx_Group_leave** (or its non-blocking form) will cause a **PMIX_GROUP_LEFT** event to be generated notifying all members of the group of the caller's departure. The function will return (or the non-blocking function will execute the specified callback function) once the event has been locally generated and is not indicative of remote receipt.

Advice to users

The `PMIx_Group_leave` API is intended solely for asynchronous departures of individual processes from a group as it is not a scalable operation – i.e., when a process determines it should no longer be a part of a defined group, but the remainder of the group retains a valid reason to continue in existence. Developers are advised to use `PMIx_Group_destruct` (or its non-blocking form) for all other scenarios as it represents a more scalable operation.

13.2.10 `PMIx_Group_leave_nb`

Summary

Non-blocking form of `PMIx_Group_leave`

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_leave_nb(const char grp[],
                   const pmix_info_t directives[], size_t ndirs,
                   pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

IN `grp`

NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group identifier (string)

IN `directives`

Array of **pmix_info_t** structures (array of handles)

IN `ndirs`

Number of elements in the *directives* array (**size_t**)

IN `cbfunc`

Callback function **pmix_op_cbfunc_t** (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed - result will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning from the API.

- 1 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
2 returned *success* - the *cbfunc* will *not* be called
- 3 • **PMIX_ERR_NOT_SUPPORTED** The PMIx library does not support this operation - the *cbfunc*
4 will *not* be called
- 5 • a PMIx error constant indicating either an error in the input or that the request was immediately
6 processed and failed - the *cbfunc* will *not* be called

7 If executed, the status returned in the provided callback function will be one of the following
8 constants:

- 9 • **PMIX_SUCCESS** The operation succeeded - i.e., the **PMIX_GROUP_LEFT** event was
10 generated
- 11 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx library supports this operation, the host RM
12 does not.
- 13 • a non-zero PMIx error constant indicating a reason for the request's failure

14  Required Attributes 
There are no identified required attributes for implementers.

15 **Description**

16 Non-blocking version of the **PMIx_Group_leave** operation. The callback function will be
17 called once the event has been locally generated and is not indicative of remote receipt.

CHAPTER 14

Network Coordinates

As the drive for performance continues, interest has grown in optimizing collective communication patterns by structuring them to follow network topology. For example, one might aggregate the contribution from all processes on a node, then again across all nodes on a common switch, and finally across all switches. Creating such optimized patterns therefore relies on detailed knowledge of the network location of each participant.

PMIx supports these efforts by defining datatypes and attributes by which network coordinates for processes and devices can be obtained from the host SMS. When used in conjunction with the PMIx *instant on* methods, this results in the ability of a process to obtain the network coordinate of all other processes without incurring additional overhead associated with the publish/exchange of that information.

14.1 Network Coordinate Datatypes

Several datatype definitions have been created to support network coordinates.

14.1.1 Network Coordinate Structure

The `pmix_coord_t` structure describes the network coordinates of a specified process in a given view

PMIx v4.0

```
typedef struct pmix_coord {
    char *fabric;
    char *plane;
    pmix_coord_view_t view;
    uint32_t *coord;
    size_t dims;
} pmix_coord_t;
```

All coordinate values shall be expressed as unsigned integers due to their units being defined in network devices and not physical distances. The coordinate is therefore an indicator of connectivity and not relative communication distance.

The fabric and plane fields are assigned by the fabric provider to help the user identify the network to which the coordinates refer. Note that providers are not required to assign any particular value to the fields and may choose to leave the fields blank. Example entries include {"Ethernet", "mgmt"} or {"infiniband", "data1"}.

Advice to PMIx library implementers

Note that the `pmix_coord_t` structure does not imply nor mandate any requirement on how the coordinate data is to be stored within the PMIx library. Implementers are free to store the coordinate in whatever format they choose.

A network coordinate is usually associated with a given network device - e.g., a particular NIC on a node. Thus, while the network coordinate of a device must be unique in a given view, the coordinate may be shared by multiple processes on a node. If the node contains multiple network devices, then either the device closest to the binding location of a process shall be used as its coordinate, or (if the process is unbound or its binding is not known) all devices on the node shall be reported as a `pmix_data_array_t` of `pmix_coord_t` structures.

Nodes with multiple network devices can also have those devices configured as multiple **network planes**. In such cases, a given process (even if bound to a specific location) may be associated with a coordinate on each plane. The resulting set of network coordinates shall be reported as a `pmix_data_array_t` of `pmix_coord_t` structures. The caller may request a coordinate from a specific network plane by passing the `PMIX_NETWORK_PLANE` attribute as a directive/qualifier to the `PMIx_Get` or `PMIx_Query_info_nb` call.

14.1.2 Network Coordinate Support Macros

The following macros are provided to support the `pmix_coord_t` structure.

14.1.2.1 Initialize the `pmix_coord_t` structure

Initialize the `pmix_coord_t` fields

PMIx v4.0

`PMIX_COORD_CONSTRUCT (m)`

IN `m`

Pointer to the structure to be initialized (pointer to `pmix_coord_t`)

14.1.2.2 Destruct the `pmix_coord_t` structure

Destruct the `pmix_coord_t` fields

PMIx v4.0

`PMIX_COORD_DESTRUCT (m)`

IN `m`

Pointer to the structure to be destructed (pointer to `pmix_coord_t`)

1 14.1.2.3 Create a `pmix_coord_t` array

2 Allocate and initialize a `pmix_coord_t` array

PMIx v4.0

C

3 `PMIX_COORD_CREATE(m, n)`

C

4 **INOUT** `m`

5 Address where the pointer to the array of `pmix_coord_t` structures shall be stored (handle)

6 **IN** `n`

7 Number of structures to be allocated (**size_t**)

8 14.1.2.4 Release a `pmix_coord_t` array

9 Release an array of `pmix_coord_t` structures

PMIx v4.0

C

10 `PMIX_COORD_FREE(m, n)`

C

11 **IN** `m`

12 Pointer to the array of `pmix_coord_t` structures (handle)

13 **IN** `n`

14 Number of structures in the array (**size_t**)

15 14.1.3 Network Coordinate Views

PMIx v4.0

C

```
16 typedef uint8_t pmix_coord_view_t;  
17 #define PMIX_COORD_VIEW_UNDEF      0x00  
18 #define PMIX_COORD_LOGICAL_VIEW    0x01  
19 #define PMIX_COORD_PHYSICAL_VIEW   0x02
```

C

20 Network coordinates can be reported based on different *views* according to user preference at the
21 time of request. The following views have been defined:

22 **PMIX_COORD_VIEW_UNDEF** The coordinate view has not been defined.

23 **PMIX_COORD_LOGICAL_VIEW** The coordinates are provided in a *logical* view, typically
24 given in Cartesian (x,y,z) dimensions, that describes the data flow in the network as defined by
25 the arrangement of the hierarchical addressing scheme, network segmentation, routing
26 domains, and other similar factors employed by that network.

27 **PMIX_COORD_PHYSICAL_VIEW** The coordinates are provided in a *physical* view based on
28 the actual wiring diagram of the network - i.e., values along each axis reflect the relative
29 position of that interface on the specific network cabling.

Advice to PMIx library implementers

PMIx library implementers are advised to avoid declaring the above constants as actual `enum` values in order to allow host environments to add support for possibly proprietary coordinate views.

If the requester does not specify a view, coordinates shall default to the *logical* view.

14.1.4 Network Coordinate Error Constants

The following error constants are used by PMIx to notify registered processes of events that affect network coordinates.

PMIX_NETWORK_COORDS_UPDATED Network coordinates have been updated - the affected networks/planes are identified in the notification. Coordinates of processes and devices on those affected components should be refreshed prior to next use.

14.1.5 Network Descriptive Attributes

These attributes are used to describe information about network resources as assigned by the RM, and thus are referenced using the process rank except where noted.

PMIX_NETWORK_COORDINATE "`pmix.net.coord`" (`pmix_data_array_t`)

Network coordinate(s) of the specified process in the view and/or plane provided by the requester. If only one NIC has been assigned to the specified process, then the array will contain only one address. Otherwise, the array will contain the coordinates of all NICs available to the process in order of least to greatest distance from the process (NICs equally distant from the process will be listed in arbitrary order).

PMIX_NETWORK_VIEW "`pmix.net.view`" (`pmix_coord_view_t`)

Network coordinate view to be used for the requested data - see `pmix_coord_view_t` for the list of accepted values.

PMIX_NETWORK_DIMS "`pmix.net.dims`" (`uint32_t`)

Request number of dimensions in the specified network plane/view. If no plane is specified, then the dimensions of all planes in the system will be returned as a `pmix_data_array_t` containing an array of `uint32_t` values. Default is to provide dimensions in *logical* view.

PMIX_NETWORK_PLANE "`pmix.net.plane`" (`char*`)

ID string of a network plane (e.g., CIDR for Ethernet). When used as a modifier in a request for information, specifies the plane whose information is to be returned. When used directly in a request, returns a `pmix_data_array_t` of string identifiers for all network planes in the system.

PMIX_NETWORK_NIC "`pmix.net.nic`" (`char*`)

1 ID string of a network interface card (NIC). When used as a modifier in a request for
2 information, specifies the NIC whose information is to be returned. When used directly in a
3 request, returns a `pmix_data_array_t` of string identifiers for all NICs in the specified
4 network plane. If no plane is specified, then the NIC identifiers of each plane in the system
5 will be returned in an array where each element is in turn an array of strings containing the
6 network plane ID followed by the identifiers of the NICs attached to that plane.

7 **PMIX_NETWORK_ENDPT** "`pmix.net.endpt`" (`pmix_data_array_t`)

8 Network endpoints for a specified process. As multiple endpoints may be assigned to a given
9 process (e.g., in the case where multiple NICs are associated with a socket to which the
10 process is bound), the returned values will be provided in a `pmix_data_array_t` - the
11 returned data type of the individual values in the array varies by fabric provider.

12 **PMIX_NETWORK_SHAPE** "`pmix.net.shape`" (`pmix_data_array_t*`)

13 The size of each dimension in the specified network plane/view, returned in a
14 `pmix_data_array_t` containing an array of `uint32_t` values. The size is defined as
15 the number of elements present in that dimension - e.g., the number of NICs in one
16 dimension of a physical view of a network plane. If no plane is specified, then the shape of
17 each plane in the system will be returned in an array of network shapes. Default is to provide
18 the shape in *logical* view.

APPENDIX A

Python Bindings

1 While the PMIx Standard is defined in terms of C-based APIs, there is no intent to limit the use of
2 PMIx to that specific language. Support for other languages is captured in the Standard by
3 describing their equivalent syntax for the PMIx APIs and native forms for the PMIx datatypes. This
4 Appendix specifically deals with Python interfaces, beginning with a review of the PMIx datatypes.
5 Support is restricted to Python 3 and above - i.e., the Python bindings do not support Python 2.

6 Note: the PMIx APIs have been loosely collected into three Python classes based on their PMIx
7 “class” (i.e., client, server, and tool). All processes have access to a basic set of the APIs, and
8 therefore those have been included in the “client” class. Servers can utilize any of those functions
9 plus a set focused on operations not commonly executed by an application process. Finally, tools
10 can also act as servers but have their own initialization function.

11 A.1 Datatype Definitions

12 PMIx defines a number of datatypes comprised of fixed-size character arrays, restricted range
13 integers (e.g., `uint32_t`), and structures. Each datatype is represented by a named unsigned 16-bit
14 integer (`uint16_t`) constant. Users are advised to use the named PMIx constants for indicating
15 datatypes instead of integer values to ensure compatibility with future PMIx versions.

16 With only a few exceptions, the C-based PMIx datatypes defined in Chapter 3 on page 19 directly
17 translate to Python. However, Python lacks the size-specific value definitions of C (e.g., `uint8_t`)
18 and thus some care must be taken to protect against overflow/underflow situations when moving
19 between the languages. Python bindings that accept values including PMIx datatypes shall
20 therefore have the datatype and associated value checked for compatibility with their PMIx-defined
21 equivalents, returning an error if:

- 22 • datatypes not defined by PMIx are encountered
- 23 • provided values fall outside the range of the C-equivalent definition - e.g., if a value identified as
24 `PMIX_UINT8` lies outside the `uint8_t` range

25 Note that explicit labeling of PMIx datatype, even when Python itself doesn’t care, is often required
26 for the Python bindings to know how to properly interpret and label the provided value when
27 passing it to the PMIx library.

28 Table A.1 lists the correspondence between datatypes in the two languages.

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>bool</code>	PMIX_BOOL	boolean	
<code>byte</code>	PMIX_BYTE	A single element byte array (i.e., a byte array of length one)	
<code>char*</code>	PMIX_STRING	string	
<code>size_t</code>	PMIX_SIZE	integer	
<code>pid_t</code>	PMIX_PID	integer	value shall be limited to the <code>uint32_t</code> range
<code>int, int8_t, int16_t, int32_t, int64_t</code>	PMIX_INT, PMIX_INT8, PMIX_INT16, PMIX_INT32, PMIX_INT64	integer	value shall be limited to its corresponding range
<code>uint, uint8_t, uint16_t, uint32_t, uint64_t</code>	PMIX_UINT, PMIX_UINT8, PMIX_UINT16, PMIX_UINT32, PMIX_UINT64	integer	value shall be limited to its corresponding range
<code>float, double</code>	PMIX_FLOAT, PMIX_DOUBLE	float	value shall be limited to its corresponding range
<code>struct timeval</code>	PMIX_TIMEVAL	{'sec': sec, 'usec': microsec}	each field is an integer value
<code>time_t</code>	PMIX_TIME	integer	limited to positive values
<code>pmix_data_type_t</code>	PMIX_DATA_TYPE	integer	value shall be limited to the <code>uint16_t</code> range
<code>pmix_status_t</code>	PMIX_STATUS	integer	
<code>pmix_key_t</code>	N/A	string	The string's length shall be limited to one less than the size of the <code>pmix_key_t</code> array (to reserve space for the terminating NULL)
<code>pmix_nspace_t</code>	N/A	string	The string's length shall be limited to one less than the size of the <code>pmix_nspace_t</code> array (to reserve space for the terminating NULL)

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_rank_t</code>	PMIX_PROC_RANK	integer	value shall be limited to the <code>uint32_t</code> range excepting the reserved values near <code>UINT32_MAX</code>
<code>pmix_proc_t</code>	PMIX_PROC	{'nspace': nspace, 'rank': rank}	<i>nspace</i> is a Python string and <i>rank</i> is an integer value. The <i>nspace</i> string's length shall be limited to one less than the size of the <code>pmix_nspace_t</code> array (to reserve space for the terminating <code>NULL</code>), and the <i>rank</i> value shall conform to the constraints associated with <code>pmix_rank_t</code>
<code>pmix_byte_object_t</code>	PMIX_BYTE_OBJECT	{'bytes': bytes, 'size': size}	<i>bytes</i> is a Python byte array and <i>size</i> is the integer number of bytes in that array.
<code>pmix_persistence_t</code>	PMIX_PERSISTENCE	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_scope_t</code>	PMIX_SCOPE	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_data_range_t</code>	PMIX_RANGE	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_proc_state_t</code>	PMIX_PROC_STATE	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_proc_info_t</code>	PMIX_PROC_INFO	{'proc': {'nspace': nspace, 'rank': rank}, 'hostname': hostname, 'executable': executable, 'pid': pid, 'exitcode': exitcode, 'state': state}	<i>proc</i> is a Python <code>proc</code> dictionary; <i>hostname</i> and <i>executable</i> are Python strings; and <i>pid</i> , <i>exitcode</i> , and <i>state</i> are Python integers

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_data_array_t</code>	PMIX_DATA_ARRAY	{'type': type, 'array': array}	<i>type</i> is the PMIx type of object in the array and <i>array</i> is a Python <i>list</i> containing the individual array elements. Note that <i>array</i> can consist of <i>any</i> PMIx types, including (for example) a Python <code>info</code> object that itself contains an <code>array</code> value
<code>pmix_info_directives_t</code>	PMIX_INFO_DIRECTIVES	integer	value shall be limited to the <code>uint32_t</code> range
<code>pmix_alloc_directive_t</code>	PMIX_ALLOC_DIRECTIVE	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_iof_channel_t</code>	PMIX_IOF_CHANNEL	integer	value shall be limited to the <code>uint16_t</code> range
<code>pmix_envar_t</code>	PMIX_ENVAR	{'envar': envar, 'value': value, 'separator': separator}	<i>envar</i> and <i>value</i> are Python strings, and <i>separator</i> a single-character Python string
<code>pmix_value_t</code>	PMIX_VALUE	{'value': value, 'val_type': type}	<i>type</i> is the PMIx datatype of <i>value</i> , and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype
<code>pmix_info_t</code>	PMIX_INFO	{'key': key, 'flags': flags, 'value': value, 'val_type': type}	<i>key</i> is a Python string <code>key</code> , <i>flags</i> is a bitmask of <code>info directives</code> , <i>type</i> is the PMIx datatype of <i>value</i> , and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype
<code>pmix_pdata_t</code>	PMIX_PDATA	{'proc': {'nspace': nspace, 'rank': rank}, 'key': key, 'value': value, 'val_type': type}	<i>proc</i> is a Python <code>proc</code> dictionary; <i>key</i> is a Python string <code>key</code> ; <i>type</i> is the PMIx datatype of <i>value</i> ; and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_app_t</code>	PMIX_APP	{'cmd': cmd, 'argv': [argv], 'env': [env], 'maxprocs': maxprocs, 'info': [info]}	<i>cmd</i> is a Python string; <i>argv</i> and <i>env</i> are Python <i>lists</i> containing Python strings; <i>maxprocs</i> is an integer; and <i>info</i> is a Python <i>list</i> of info values
<code>pmix_query_t</code>	PMIX_QUERY	{'keys': [keys], 'qualifiers': [info]}	<i>keys</i> is a Python <i>list</i> of Python strings, and <i>qualifiers</i> is a Python <i>list</i> of info values
<code>pmix_regattr_t</code>	PMIX_REGATTR	{'name': name, 'key': key, 'type': type, 'info': [info], 'description': [desc]}	<i>name</i> and <i>string</i> are Python strings; <i>type</i> is the PMIx datatype for the attribute's value; <i>info</i> is a Python <i>list</i> of info values; and <i>description</i> is a list of Python strings describing the attribute

1 A.1.1 Example

2 Converting a C-based program to its Python equivalent requires translation of the relevant
3 datatypes as well as use of the appropriate API form. An example small program may help
4 illustrate the changes. Consider the following C-based program snippet:

```
5 #include <pmix.h>  
6 ...  
7  
8 pmix_info_t info[2];  
9  
10 PMIX_INFO_LOAD(&info[0], PMIX_PROGRAMMING_MODEL, "TEST", PMIX_STRING)  
11 PMIX_INFO_LOAD(&info[1], PMIX_MODEL_LIBRARY_NAME, "PMIX", PMIX_STRING)  
12  
13 rc = PMIx_Init(&myproc, info, 2);  
14  
15 PMIX_INFO_DESTRUCT(&info[0]); // free the copied string  
16 PMIX_INFO_DESTRUCT(&info[1]); // free the copied string  
17
```

18 Moving to the Python version requires that the `pmix_info_t` be translated to the Python `info`
19 equivalent, and that the returned information be captured in the return parameters as opposed to a
20 pointer parameter in the function call, as shown below:

```
21 import pmix  
22 ...  
23  
24 myclient = PMIxClient()  
25 info = [{'key':PMIX_PROGRAMMING_MODEL,  
26         'value':'TEST', 'val_type':PMIX_STRING},  
27         {'key':PMIX_MODEL_LIBRARY_NAME,  
28         'value':'PMIX', 'val_type':PMIX_STRING}]  
29 rc,myproc = myclient.init(info)  
30
```

31 Note the use of the `PMIX_STRING` identifier to ensure the Python bindings interpret the provided
32 string value as a PMIx "string" and not an array of bytes.

33 A.2 Function Definitions

34 A.2.1 IOF Delivery Function

35 Summary

36 Callback function for delivering forwarded IO to a process

1
PMIx v4.0

Format

Python

```
2 def iofcbfunc(iofhdlr:integer, channel:integer,  
3             source:dict, payload:dict, info:list)
```

Python

4 **IN** **iofhdlr**

Registration number of the handler being invoked (integer)

6 **IN** **channel**

Python **channel** bitmask identifying the channel the data arrived on (integer)

8 **IN** **source**

Python **proc** identifying the namespace/rank of the process that generated the data (dict)

10 **IN** **payload**

Python **byteobject** containing the data (dict)

12 **IN** **info**

List of Python **info** provided by the source containing metadata about the payload. This could include **PMIX_IOF_COMPLETE** (list)

15 Returns: nothing

16 See [pmix_iof_cbfunc_t](#) for details

17 A.2.2 Event Handler

18 Summary

19 Callback function for event handlers

20 Format

PMIx v4.0

Python

```
21 def evhandler(evhdlr:integer, status:integer,  
22             source:dict, info:list, results:list)
```

Python

23 **IN** **iofhdlr**

Registration number of the handler being invoked (integer)

25 **IN** **status**

Status associated with the operation (integer)

27 **IN** **source**

Python **proc** identifying the namespace/rank of the process that generated the event (dict)

29 **IN** **info**

List of Python **info** provided by the source containing metadata about the event (list)

31 **IN** **results**

List of Python **info** containing the aggregated results of all prior evhandlers (list)

33 Returns:

- 1 • *rc* - Status returned by the event handler's operation (integer)
- 2 • *results* - List of Python **info** containing results from this event handler's operation on the event
- 3 (list)
- 4 See [pmix_notification_fn_t](#) for details

5 A.2.3 Server Module Functions

6 The following definitions represent functions that may be provided to the PMIx server library at
7 time of initialization for servicing of client requests. Module functions that are not provided default
8 to returning "not supported" to the caller.



9 A.2.3.1 Client Connected

10 Summary

11 Notify the host server that a client connected to this server.

12 Format

PMIx v4.0

```
Python   
def clientconnected(proc:dict is not None)  
Python 
```

14 IN **proc**

15 Python **proc** identifying the namespace/rank of the process that connected (dict)

16 Returns:

- 17 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the connection should be rejected
- 18 (integer)

19 See [pmix_server_client_connected_fn_t](#) for details



20 A.2.3.2 Client Finalized

21 Summary

22 Notify the host environment that a client called **PMIx_Finalize**.

23 Format

PMIx v4.0

```
Python   
def clientfinalized(proc:dict is not None):  
Python 
```

25 IN **proc**

26 Python **proc** identifying the namespace/rank of the process that finalized (dict)

27 Returns: nothing

28 See [pmix_server_client_finalized_fn_t](#) for details

1 **A.2.3.3 Client Aborted**

2 **Summary**

3 Notify the host environment that a local client called `PMIx_Abort` .

4 **Format**

PMIx v4.0

Python

```
5 def clientaborted(proc:dict is not None, status:integer,  
6                 msg:str, targets:list)
```

Python

7 **IN** `proc`

8 Python `proc` identifying the namespace/rank of the process that called abort (dict)

9 **IN** `status`

10 PMIx status to be returned on exit (integer)

11 **IN** `msg`

12 String message to be printed (string)

13 **IN** `targets`

14 List of Python `proc` dictionaries (list)

15 Returns:

- 16 • `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)

17 See `pmix_server_abort_fn_t` for details

18 **A.2.3.4 Fence**

19 **Summary**

20 At least one client called either `PMIx_Fence` or `PMIx_Fence_nb`

21 **Format**

PMIx v4.0

Python

```
22 def fence(procs:list, directives:list, data:bytearray)
```

Python

23 **IN** `procs`

24 List of Python `proc` dictionaries (list)

25 **IN** `directives`

26 List of Python `info` dictionaries (list)

27 **IN** `data`

28 Python bytearray of data to be circulated during fence operation (bytearray)

29 Returns:

- 30 • `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)
- 31 • `data` - Python bytearray containing the aggregated data from all participants (bytearray)

32 See `pmix_server_fence_nb_fn_t` for details

1 **A.2.3.5 Direct Modex**

2 **Summary**

3 Used by the PMIx server to request its local host contact the PMIx server on the remote node that
4 hosts the specified proc to obtain and return a direct modex blob for that proc.

5 **Format**

PMIx v4.0

▼ Python ▼

6 **def dmodex(proc:dict, directives:list)**

▲ Python ▲

7 **IN proc**

8 Python **proc** dictionary of process whose data is being requested (list)

9 **IN directives**

10 List of Python **info** dictionaries (list)

11 Returns:

12 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

13 • *data* - Python bytearray containing the data for the specified process (bytearray)

14 See [pmix_server_dmodex_req_fn_t](#) for details

15 **A.2.3.6 Publish**

16 **Summary**

17 Publish data per the PMIx API specification.

18 **Format**

PMIx v4.0

▼ Python ▼

19 **def publish(proc:dict, directives:list)**

▲ Python ▲

20 **IN proc**

21 Python **proc** dictionary of process publishing the data (list)

22 **IN directives**

23 List of Python **info** dictionaries containing data and directives (list)

24 Returns:

25 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

26 See [pmix_server_publish_fn_t](#) for details

27 **A.2.3.7 Lookup**

28 **Summary**

29 Lookup published data.

1 **Format**
PMIx v4.0 Python

2 `def lookup(proc:dict, keys:list, directives:list)`
Python

- 3 **IN proc**
4 Python `proc` dictionary of process seeking the data (list)
- 5 **IN keys**
6 List of Python strings (list)
- 7 **IN directives**
8 List of Python `info` dictionaries containing directives (list)

9 Returns:
10

- `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)
- `pdata` - List of `pdata` containing the returned results (list)

11 See `pmix_server_lookup_fn_t` for details

13 A.2.3.8 Unpublish

14 **Summary**
15 Delete data from the data store.

16 **Format**
PMIx v4.0 Python

17 `def unpublish(proc:dict, keys:list, directives:list)`
Python

- 18 **IN proc**
19 Python `proc` dictionary of process making the request (list)
- 20 **IN keys**
21 List of Python strings (list)
- 22 **IN directives**
23 List of Python `info` dictionaries containing directives (list)

24 Returns:
25

- `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)

26 See `pmix_server_unpublish_fn_t` for details

27 A.2.3.9 Spawn

28 **Summary**
29 Spawn a set of applications/processes as per the `PMIx_Spawn` API.

1
PMIx v4.0

Format

Python

2 **def spawn(proc:dict, jobInfo:list, apps:list)**

Python

3 **IN proc**

Python **proc** dictionary of process making the request (list)

4 **IN jobInfo**

List of Python **info** job-level directives and information (list)

5 **IN apps**

List of Python **app** dictionaries describing applications to be spawned (list)

6 Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- *nspac* - Python string containing namespace of the spawned job (str)

See [pmix_server_spawn_fn_t](#) for details

A.2.3.10 Connect

Summary

Record the specified processes as *connected*.

PMIx v4.0

Format

Python

7 **def connect(procs:list, directives:list)**

Python

8 **IN procs**

List of Python **proc** dictionaries identifying participants (list)

9 **IN directives**

List of Python **info** directives (list)

10 Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

See [pmix_server_connect_fn_t](#) for details

A.2.3.11 Disconnect

Summary

Disconnect a previously connected set of processes.

1 **Format**

PMIx v4.0 Python

2 `def disconnect(procs:list, directives:list)`
Python

3 **IN procs**
List of Python **proc** dictionaries identifying participants (list)

4 **IN directives**
List of Python **info** directives (list)

5

6

7 Returns:

8 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

9 See [pmix_server_disconnect_fn_t](#) for details

10 A.2.3.12 Register Events

11 **Summary**
Register to receive notifications for the specified events.

13 **Format**

PMIx v4.0 Python

14 `def register_events(codes:list, directives:list)`
Python

15 **IN codes**
List of Python integers (list)

16 **IN directives**
List of Python **info** directives (list)

17

18

19 Returns:

20 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

21 See [pmix_server_register_events_fn_t](#) for details

22 A.2.3.13 Deregister Events

23 **Summary**
Deregister to receive notifications for the specified events.

24

1 **Format**
PMIx v4.0 Python

2 **def deregister_events (codes:list)**
Python

3 **IN codes**
4 List of Python integers (list)

5 Returns:

- 6 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

7 See [pmix_server_deregister_events_fn_t](#) for details

8 **A.2.3.14 Notify Event**

9 **Summary**

10 Notify the specified range of processes of an event.

11 **Format**
PMIx v4.0 Python

12 **def notify_event (code:integer, source:dict, range:integer, directives:list)**
Python

13 **IN code**
14 Python integer [pmix_status_t](#) (list)

15 **IN source**
16 Python [proc](#) of process that generated the event (dict)

17 **IN range**
18 Python [range](#) in which the event is to be reported (integer)

19 **IN directives**
20 List of Python [info](#) directives (list)

21 Returns:

- 22 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

23 See [pmix_server_notify_event_fn_t](#) for details

24 **A.2.3.15 Query**

25 **Summary**

26 Query information from the resource manager.

1 **Format**
PMIx v4.0 Python

2 **def query(proc:dict, queries:list)**
Python

3 **IN proc**
4 Python **proc** of requesting process (dict)
5 **IN queries**
6 List of Python **query** directives (list)

7 Returns:

- 8 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 9 • *info* - List of Python **info** containing the returned results (list)

10 See [pmix_server_query_fn_t](#) for details

11 A.2.3.16 Tool Connected

12 Summary

13 Register that a tool has connected to the server.

14 **Format**
PMIx v4.0 Python

15 **def tool_connected(info:list)**
Python

16 **IN info**
17 List of Python **info** containing info on the connecting tool (list)

18 Returns:


- 19 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 20 • *proc* - Python **proc** containing the assigned namespace:rank for the tool (dict)


21 See [pmix_server_tool_connection_fn_t](#) for details

22 A.2.3.17 Log

23 Summary

24 Log data on behalf of a client.

1 **Format** Python 
PMIx v4.0

2 **def log(proc:dict, data:list, directives:list)** Python 

3 **IN proc**
 4 Python **proc** of requesting process (dict)

5 **IN data**
 6 List of Python **info** containing data to be logged (list)

7 **IN directives**
 8 List of Python **info** containing directives (list)

9 Returns:

10 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

11 See [pmix_server_log_fn_t](#) for details

12 **A.2.3.18 Allocate Resources**

13 **Summary**

14 Request allocation operations on behalf of a client.

15 **Format** Python

PMIx v4.0

16 **def allocate(proc:dict, action:integer, directives:list)** Python 

17 **IN proc**
 18 Python **proc** of requesting process (dict)

19 **IN action**
 20 Python **allocdir** specifying requested action (integer)

21 **IN directives**
 22 List of Python **info** containing directives (list)

23 Returns:

24 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

25 • *refarginfo* - List of Python **info** containing results of requested operation (list)

26 See [pmix_server_alloc_fn_t](#) for details

27 **A.2.3.19 Job Control**

28 **Summary**

29 Execute a job control action on behalf of a client.

1
PMIx v4.0

Format

Python

2 `def job_control(proc:dict, targets:list, directives:list)`

Python

3 **IN** `proc`

Python `proc` of requesting process (dict)

4 **IN** `targets`

List of Python `proc` specifying target processes (list)

5 **IN** `directives`

List of Python `info` containing directives (list)

6 Returns:

- `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)

See [pmix_server_job_control_fn_t](#) for details

A.2.3.20 Monitor

Summary

Request that a client be monitored for activity.

Format

Python

PMIx v4.0

16 `def monitor(proc:dict, request:list, directives:list)`

Python

17 **IN** `proc`

Python `proc` of requesting process (dict)

18 **IN** `request`

List of Python `info` specifying requested monitoring operations (list)

19 **IN** `directives`

List of Python `info` containing directives (list)

20 Returns:

- `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the operation failed (integer)

See [pmix_server_monitor_fn_t](#) for details

A.2.3.21 Get Credential

Summary

Request a credential from the host environment

1 **Format**
PMIx v4.0 Python

2 `def get_credential(proc:dict, directives:list)`
Python

3 **IN proc**
4 Python **proc** of requesting process (dict)

5 **IN directives**
6 List of Python **info** containing directives (list)

7 Returns:

- 8 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 9 • *cred* - Python **byteobject** containing returned credential (dict)
- 10 • *info* - List of Python **info** containing any additional info about the credential (list)

11 See [pmix_server_get_cred_fn_t](#) for details

12 **A.2.3.22 Validate Credential**

13 **Summary**

14 Request validation of a credential

15 **Format**
PMIx v4.0 Python

16 `def validate_credential(proc:dict, cred:dict, directives:list)`
Python

17 **IN proc**
18 Python **proc** of requesting process (dict)

19 **IN cred**
20 Python **byteobject** containing credential (dict)

21 **IN directives**
22 List of Python **info** containing directives (list)

23 Returns:

- 24 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 25 • *info* - List of Python **info** containing any additional info from the credential (list)

26 See [pmix_server_validate_cred_fn_t](#) for details

27 **A.2.3.23 IO Forward**

28 **Summary**

29 Request the specified IO channels be forwarded from the given array of processes.

1 **Format**

PMIx v4.0 Python

```
2 def iof_pull(sources:list, channels:integer, directives:list)
```

Python

3 **IN sources**
 4 List of Python [proc](#) whose IO is being requested (list)

5 **IN channels**
 6 Bitmask of Python [channel](#) identifying IO channels to be forwarded (integer)

7 **IN directives**
 8 List of Python [info](#) containing directives (list)

9 Returns:

10 • *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)

11 See [pmix_server_iof_fn_t](#) for details

12 A.2.3.24 IO Push

13 Summary

14 Pass standard input data to the host environment for transmission to specified recipients.

15 Format

PMIx v4.0 Python

```
16 def iof_push(source:dict, targets:list, directives:list)
```

Python

17 **IN source**
 18 Python [proc](#) whose stdin data is being provided (dict)

19 **IN targets**
 20 List of Python [proc](#) identifying targets to receive the provided data (list)

21 **IN directives**
 22 List of Python [info](#) containing directives (list)

23 Returns:

24 • *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)

25 See [pmix_server_stdin_fn_t](#) for details

26 A.3 PMIxClient

27 The client Python class is by far the richest in terms of APIs as it houses all the APIs that an
 28 application might utilize. Due to the datatype translation requirements of the C-Python interface,
 29 only the blocking form of each API is supported – providing a Python callback function directly to
 30 the C interface underlying the bindings was not a supportable option.

1 **A.3.1 Client.init**

2 **Summary**

3 Initialize the PMIx client library after obtaining a new PMIxCliient object

4 **Format**

PMIx v4.0

Python 

5 `rc, proc = myclient.init(info:list)`

Python 

6 **IN** `info`

7 List of Python `info` dictionaries (list)

8 Returns:

- 9 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 10 • `proc` - a Python `proc` dictionary (dict)

11 See **PMIx_Init** for description of all relevant attributes and behaviors

12 **A.3.2 Client.initialized**

13 **Format**

PMIx v4.0

Python 

14 `rc = myclient.initialized()`

Python 

15 Returns:

- 16 • `rc` - a value of **1** (true) will be returned if the PMIx library has been initialized, and **0** (false) otherwise (integer)

18 See **PMIx_Initialized** for description of all relevant attributes and behaviors

19 **A.3.3 Client.get_version**

20 **Format**

PMIx v4.0

Python 

21 `vers = myclient.get_version()`

Python 

22 Returns:

- 23 • `vers` - Python string containing the version of the PMIx library (e.g., "3.1.4") (integer)

24 See **PMIx_Get_version** for description of all relevant attributes and behaviors

1 A.3.4 Client.finalize

2 Summary

3 Finalize the PMIx client library.

4 Format

PMIx v4.0

▼ Python ▼

5 `rc = myclient.finalize(info:list)`

▲ Python ▲

6 **IN info**

7 List of Python **info** dictionaries (list)

8 Returns:

- 9 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

10 See **PMIx_Finalize** for description of all relevant attributes and behaviors

11 A.3.5 Client.abort

12 Summary

13 Request that the provided list of procs be aborted

14 Format

PMIx v4.0

▼ Python ▼

15 `rc = myclient.abort(status:integer, msg:str, targets:list)`

▲ Python ▲

16 **IN status**

17 PMIx status to be returned on exit (integer)

18 **IN msg**

19 String message to be printed (string)

20 **IN targets**

21 List of Python **proc** dictionaries (list)

22 Returns:

- 23 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

24 See **PMIx_Abort** for description of all relevant attributes and behaviors

25 A.3.6 Client.store_internal

26 Summary

27 Store some data locally for retrieval by other areas of the process

1 **Format**

PMIx v4.0 Python

2 `rc = myclient.store_internal(proc:dict, key:str, value:dict)`
Python

3 **IN proc**
4 Python `proc` dictionary of the process being referenced (dict)

5 **IN key**
6 String key of the data (string)

7 **IN value**
8 Python `value` dictionary (dict)

9 Returns:

10 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

11 See `PMIx_Store_internal` for details

12 A.3.7 Client.put

13 **Summary**

14 Push a key/value pair into the client's namespace.

15 **Format**

PMIx v4.0 Python

16 `rc = myclient.put(scope:integer, key:str, value:dict)`
Python

17 **IN scope**
18 Scope of the data being posted (integer)

19 **IN key**
20 String key of the data (string)

21 **IN value**
22 Python `value` dictionary (dict)

23 Returns:

24 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

25 See `PMIx_Put` for description of all relevant attributes and behaviors

26 A.3.8 Client.commit

27 **Summary**

28 Push all previously `PMIxClient.put` values to the local PMIx server.

1 **Format**
PMIx v4.0 Python

2 `rc = myclient.commit()` Python

3 Returns:

- 4 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

5 See **PMIx_Commit** for description of all relevant attributes and behaviors

6 **A.3.9 Client.fence**

7 **Summary**

8 Execute a blocking barrier across the processes identified in the specified list

9 **Format**
PMIx v4.0 Python

10 `rc = myclient.fence(peers:list, directives:list)` Python

11 **IN** `peers`
12 List of Python **proc** dictionaries (list)

13 **IN** `directives`
14 List of Python **info** dictionaries (list)

15 Returns:

- 16 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

17 See **PMIx_Fence** for description of all relevant attributes and behaviors

18 **A.3.10 Client.get**

19 **Summary**

20 Retrieve a key/value pair

1
PMIx v4.0

Format

Python

2 `rc, val = myclient.get(proc=dict, key=str, directives=list)`

Python

- 3 **IN** `proc`
4 Python `proc` whose data is being requested (dict)
5 **IN** `key`
6 Python string key of the data to be returned (str)
7 **IN** `directives`
8 List of Python `info` dictionaries (list)

9 Returns:

- 10 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
11 • `val` - Python `value` containing the returned data (dict)

12 See `PMIx_Get` for description of all relevant attributes and behaviors

13 A.3.11 Client.publish

14 Summary

15 Publish data for later access via `PMIx_Lookup` .

16 Format

Python

PMIx v4.0

17 `rc = myclient.publish(directives=list)`

Python

- 18 **IN** `directives`
19 List of Python `info` dictionaries containing data to be published and directives (list)

20 Returns:

- 21 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

22 See `PMIx_Publish` for description of all relevant attributes and behaviors

23 A.3.12 Client.lookup

24 Summary

25 Lookup information published by this or another process with `PMIx_Publish` .

1 **Format**
PMIx v4.0 Python

```
2 rc, info = myclient.lookup(pdata:list, directives:list)
```

Python

3 **IN** **pdata**
 4 List of Python **pdata** dictionaries identifying data to be retrieved (list)

5 **IN** **directives**
 6 List of Python **info** dictionaries (list)

7 Returns:

- 8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 9 • *info* - Python list of **info** containing the returned data (list)

10 See **PMIx_Lookup** for description of all relevant attributes and behaviors

11 A.3.13 Client.unpublish

12 **Summary**
 13 Delete data published by this process with **PMIx_Publish** .

14 **Format**
PMIx v4.0 Python

```
15 rc = myclient.unpublish(keys:list, directives:list)
```

Python

16 **IN** **keys**
 17 List of Python string keys identifying data to be deleted (list)

18 **IN** **directives**
 19 List of Python **info** dictionaries (list)

20 Returns:

- 21 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

22 See **PMIx_Unpublish** for description of all relevant attributes and behaviors

23 A.3.14 Client.spawn

24 **Summary**
 25 Spawn a new job.

1 **Format**
PMIx v4.0 Python

```
2 rc, nspace = myclient.spawn(jobinfo:list, apps:list)
```

3 **IN** **jobinfo**
4 List of Python **info** dictionaries (list)

5 **IN** **apps**
6 List of Python **app** dictionaries (list)

7 Returns:

- 8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 9 • *nspace* - Python **nspace** of the new job (dict)

10 See **PMIx_Spawn** for description of all relevant attributes and behaviors

11 A.3.15 Client.connect

12 **Summary**
13 Connect namespaces.

14 **Format**
PMIx v4.0 Python

```
15 rc = myclient.connect(peers:list, directives:list)
```

16 **IN** **peers**
17 List of Python **proc** dictionaries (list)

18 **IN** **directives**
19 List of Python **info** dictionaries (list)

20 Returns:

- 21 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

22 See **PMIx_Connect** for description of all relevant attributes and behaviors

23 A.3.16 Client.disconnect

24 **Summary**
25 Disconnect namespaces.

1 **Format**
PMIx v4.0 Python

2 `rc = myclient.disconnect(peers:list, directives:list)`
 Python

3 **IN peers**
 4 List of Python **proc** dictionaries (list)

5 **IN directives**
 6 List of Python **info** dictionaries (list)

7 Returns:

8 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 See **PMIx_Disconnect** for description of all relevant attributes and behaviors

10 A.3.17 Client.resolve_peers

11 **Summary**
 12 Return list of processes within the specified **nospace** on the given node.

13 **Format**
PMIx v4.0 Python

14 `rc,procs = myclient.resolve_peers(node:str, nospace:str)`
 Python

15 **IN node**
 16 Name of node whose processes are being requested (str)

17 **IN nospace**
 18 Python **nospace** whose processes are to be returned (str)

19 Returns:

20 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

21 • `procs` - List of Python **proc** dictionaries (list)

22 See **PMIx_Resolve_peers** for description of all relevant attributes and behaviors

23 A.3.18 Client.resolve_nodes

24 **Summary**
 25 Return list of nodes hosting processes within the specified **nospace** .

1 **Format**
PMIx v4.0 Python

```
2 rc, nodes = myclient.resolve_nodes(namespace: str)
```

3 **IN namespace**
4 Python **namespace** (str)

5 Returns:

- 6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 7 • *nodes* - List of Python string node names (list)

8 See **PMIx_Resolve_nodes** for description of all relevant attributes and behaviors

9 **A.3.19 Client.query**

10 **Summary**
11 Query information about the system in general

12 **Format**
PMIx v4.0 Python

```
13 rc, info = myclient.query(queries: list, directives: list)
```

14 Python

14 **IN queries**
15 List of Python **query** dictionaries (list)

16 **IN directives**
17 List of Python **info** dictionaries (list)

18 Returns:

- 19 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 20 • *info* - List of Python **info** containing results of the query (list)

21 See **PMIx_Query_info_nb** for description of all relevant attributes and behaviors

22 **A.3.20 Client.log**

23 **Summary**
24 Log data to a central data service/store

1 **Format**
PMIx v4.0 Python

2 `rc = myclient.log(data:list, directives:list)`
 Python

3 **IN data**
 4 List of Python **info** dictionaries (list)

5 **IN directives**
 6 List of Python **info** dictionaries (list)

7 Returns:

8 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 See **PMIx_Log** for description of all relevant attributes and behaviors

10 A.3.21 Client.allocate

11 **Summary**
 12 Request an allocation operation from the host resource manager.

13 **Format**
PMIx v4.0 Python

14 `rc,info = myclient.allocate(request:integer, directives:list)`
 Python

15 **IN request**
 16 Python **allocdir** specifying requested operation (integer)

17 **IN directives**
 18 List of Python **info** dictionaries describing request (list)

19 Returns:

20 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

21 • `info` - List of Python **info** containing results of the request (list)

22 See **PMIx_Allocation_request_nb** for description of all relevant attributes and behaviors

23 A.3.22 Client.job_ctrl

24 **Summary**
 25 Request a job control action

1 **Format**

PMIx v4.0 Python

```
2 rc,info = myclient.job_ctrl(targets:list, directives:list)
```

Python

3 **IN targets**

4 List of Python **proc** specifying targets of requested operation (integer)

5 **IN directives**

6 List of Python **info** dictionaries describing request (list)

7 Returns:

8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 • *info* - List of Python **info** containing results of the request (list)

10 See **PMIx_Job_control_nb** for description of all relevant attributes and behaviors

11 A.3.23 Client.monitor

12 **Summary**

13 Request that something be monitored

14 **Format**

PMIx v4.0 Python

```
15 rc,info = myclient.monitor(targets:list, directives:list)
```

Python

16 **IN targets**

17 List of Python **proc** specifying targets of requested operation (integer)

18 **IN directives**

19 List of Python **info** dictionaries describing request (list)

20 Returns:

21 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

22 • *info* - List of Python **info** containing results of the request (list)

23 See **PMIx_Process_monitor_nb** for description of all relevant attributes and behaviors

24 A.3.24 Client.get_credential

25 **Summary**

26 Request a credential from the PMIx server/SMS

1 **Format**

PMIx v4.0 Python

```
2 rc, cred, info = myclient.get_credential(directives:list)
```

Python

3 **IN directives**

4 List of Python **info** dictionaries describing request (list)

5 Returns:

6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

7 • *cred* - Python **byteobject** containing returned credential (dict)

8 • *info* - List of Python **info** containing results of the request (list)

9 See **PMIx_Get_credential** for description of all relevant attributes and behaviors

10 A.3.25 Client.validate_credential

11 **Summary**

12 Request validation of a credential by the PMIx server/SMS

13 **Format**

PMIx v4.0 Python

```
14 rc, info = myclient.validate_credential(cred:dict, directives:list)
```

Python

15 **IN cred**

16 Python **byteobject** containing credential (dict)

17 **IN directives**

18 List of Python **info** dictionaries describing request (list)

19 Returns:

20 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

21 • *info* - List of Python **info** containing additional results of the request (list)

22 See **PMIx_Validate_credential** for description of all relevant attributes and behaviors

23 A.3.26 Client.group_construct

24 **Summary**

25 Construct a new group composed of the specified processes and identified with the provided group

26 identifier

1 **Format**

2 *PMIx v4.0* `rc, info = myclient.construct_group(grp:string, members:list, directives:list)` Python

3 **IN grp**
 4 Python string identifier for the group (str)

5 **IN members**
 6 List of Python **proc** dictionaries identifying group members (list)

7 **IN directives**
 8 List of Python **info** dictionaries describing request (list)

9 Returns:

10 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

11 • *info* - List of Python **info** containing results of the request (list)

12 See **PMIx_Group_construct** for description of all relevant attributes and behaviors

13 A.3.27 Client.group_invite

14 **Summary**
 15 Explicitly invite specified processes to join a group

16 **Format**

17 *PMIx v4.0* `rc, info = myclient.group_invite(grp:string, members:list, directives:list)` Python

18 **IN grp**
 19 Python string identifier for the group (str)

20 **IN members**
 21 List of Python **proc** dictionaries identifying processes to be invited (list)

22 **IN directives**
 23 List of Python **info** dictionaries describing request (list)

24 Returns:

25 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

26 • *info* - List of Python **info** containing results of the request (list)

27 See **PMIx_Group_invite** for description of all relevant attributes and behaviors

1 A.3.28 Client.group_join

2 Summary

3 Respond to an invitation to join a group that is being asynchronously constructed

4 Format

PMIx v4.0

```
Python
```

```
rc, info = myclient.group_join(grp:string, leader:dict, opt:integer, directives: list)
```

```
Python
```

6 IN grp

7 Python string identifier for the group (str)

8 IN leader

9 Python [proc](#) dictionary identifying process leading the group (dict)

10 IN opt

11 One of the [pmix_group_opt_t](#) values indicating decline/accept (integer)

12 IN directives

13 List of Python [info](#) dictionaries describing request (list)

14 Returns:

- 15 • *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- 16 • *info* - List of Python [info](#) containing results of the request (list)

17 See [PMIx_Group_join](#) for description of all relevant attributes and behaviors

18 A.3.29 Client.group_leave

19 Summary

20 Leave a PMIx Group

21 Format

PMIx v4.0

```
Python
```

```
rc = myclient.group_leave(grp:string, directives:list)
```

```
Python
```

23 IN grp

24 Python string identifier for the group (str)

25 IN directives

26 List of Python [info](#) dictionaries describing request (list)

27 Returns:

- 28 • *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

29 See [PMIx_Group_leave](#) for description of all relevant attributes and behaviors

1 A.3.30 Client.group_destruct

2 Summary

3 Destruct a PMIx Group

4 Format

PMIx v4.0

Python

```
5 rc = myclient.group_destruct(grp:string, directives:list)
```

Python

6 **IN grp**

7 Python string identifier for the group (str)

8 **IN directives**

9 List of Python **info** dictionaries describing request (list)

10 Returns:

- 11 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

12 See **PMIx_Group_destruct** for description of all relevant attributes and behaviors

13 A.3.31 Client.register_event_handler

14 Summary

15 Register an event handler to report events.

16 Format

PMIx v4.0

Python

```
17 rc,id = myclient.register_event_handler(codes:list, directives:list, cbfunc)
```

Python

18 **IN codes**

19 List of Python integer status codes that should be reported to this handler (llist)

20 **IN directives**

21 List of Python **info** dictionaries describing request (list)

22 **IN cbfunc**

23 Python **evhandler** to be called when event is received (func)

24 Returns:

- 25 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

- 26 • *id* - PMIx reference identifier for handler (integer)

27 See **PMIx_Register_event_handler** for description of all relevant attributes and behaviors

1 A.3.32 Client.deregister_event_handler

2 Summary

3 Deregister an event handler

4 Format

PMIx v4.0

Python

5 `myclient.deregister_event_handler(id:integer)`

Python

6 **IN** `id`

7 PMIx reference identifier for handler (integer)

8 Returns: None

9 See [PMIx_Deregister_event_handler](#) for description of all relevant attributes and
10 behaviors

11 A.3.33 Client.notify_event

12 Summary

13 Report an event for notification via any registered handler.

14 Format

PMIx v4.0

Python

15 `rc = myclient.notify_event(status:integer, source:dict,`
16 `range:integer, directives:list)`

Python

17 **IN** `status`

18 PMIx status code indicating the event being reported (integer)

19 **IN** `source`

20 Python [proc](#) of the process that generated the event (dict)

21 **IN** `range`

22 Python [range](#) in which the event is to be reported (integer)

23 **IN** `directives`

24 List of Python [info](#) dictionaries describing request (list)

25 Returns:

26 • `rc` - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

27 See [PMIx_Notify_event](#) for description of all relevant attributes and behaviors

1 **A.3.34 Client.error_string**

2 **Summary**

3 Pretty-print string representation of [pmix_status_t](#).

4 **Format**

PMIx v4.0

▼ Python ▼

5 `rep = myclient.error_string(status:integer)`

▲ Python ▲

6 **IN status**

7 PMIx status code (integer)

8 Returns:

- 9 • *rep* - String representation of the provided status code (str)

10 See [PMIx_Error_string](#) for further details

11 **A.3.35 Client.proc_state_string**

12 **Summary**

13 Pretty-print string representation of [pmix_proc_state_t](#).

14 **Format**

PMIx v4.0

▼ Python ▼

15 `rep = myclient.proc_state_string(state:integer)`

▲ Python ▲

16 **IN state**

17 PMIx process state code (integer)

18 Returns:

- 19 • *rep* - String representation of the provided process state (str)

20 See [PMIx_Proc_state_string](#) for further details

21 **A.3.36 Client.scope_string**

22 **Summary**

23 Pretty-print string representation of [pmix_scope_t](#).

1 **Format**
PMIx v4.0 Python

2 `rep = myclient.scope_string(scope:integer)`
Python

3 **IN** `scope`
4 PMIx scope value (integer)

5 Returns:

- 6 • `rep` - String representation of the provided scope (str)

7 See [PMIx Scope string](#) for further details

8 **A.3.37 Client.persistence_string**

9 **Summary**

10 Pretty-print string representation of [pmix_persistence_t](#).

11 **Format**
PMIx v4.0 Python

12 `rep = myclient.persistence_string(persistence:integer)`
Python

13 **IN** `persistence`
14 PMIx persistence value (integer)

15 Returns:

- 16 • `rep` - String representation of the provided persistence (str)

17 See [PMIx Persistence string](#) for further details

18 **A.3.38 Client.data_range_string**

19 **Summary**

20 Pretty-print string representation of [pmix_data_range_t](#).

21 **Format**
PMIx v4.0 Python

22 `rep = myclient.data_range_string(range:integer)`
Python

23 **IN** `range`
24 PMIx data range value (integer)

25 Returns:

- 26 • `rep` - String representation of the provided data range (str)

27 See [PMIx Data range string](#) for further details

1 A.3.39 Client.info_directives_string

2 Summary

3 Pretty-print string representation of [pmix_info_directives_t](#) .

4 Format

PMIx v4.0

▼ Python ▼

5 `rep = myclient.info_directives_string(directives:integer)`

▲ Python ▲

6 IN directives

7 PMIx info directives value (integer)

8 Returns:

- 9 • *rep* - String representation of the provided info directives (str)

10 See [PMIx_Info_directives_string](#) for further details

11 A.3.40 Client.data_type_string

12 Summary

13 Pretty-print string representation of [pmix_data_type_t](#) .

14 Format

PMIx v4.0

▼ Python ▼

15 `rep = myclient.data_type_string(dtype:integer)`

▲ Python ▲

16 IN dtype

17 PMIx datatype value (integer)

18 Returns:

- 19 • *rep* - String representation of the provided datatype (str)

20 See [PMIx_Data_type_string](#) for further details

21 A.3.41 Client.alloc_directive_string

22 Summary

23 Pretty-print string representation of [pmix_alloc_directive_t](#) .

1 **Format**
PMIx v4.0 Python

```
2 rep = myclient.alloc_directive_string(adir:integer)
```

3 **IN** `adir`
4 PMIx allocation directive value (integer)

5 Returns:

- 6 • `rep` - String representation of the provided allocation directive (str)

7 See [PMIx_Alloc_directive_string](#) for further details

8 **A.3.42 Client.iof_channel_string**

9 **Summary**
10 Pretty-print string representation of `pmix_iof_channel_t`.

11 **Format**
PMIx v4.0 Python

```
12 rep = myclient.iof_channel_string(channel:integer)
```

13 **IN** `channel`
14 PMIx IOF channel value (integer)

15 Returns:

- 16 • `rep` - String representation of the provided IOF channel (str)

17 See [PMIx_IOF_channel_string](#) for further details

18 **A.4 PMIxServer**

19 The server Python class inherits the Python "client" class as its parent. Thus, it includes all client
20 functions in addition to the ones defined in this section.

21 **A.4.1 Server.init**

22 **Summary**
23 Initialize the PMIx server library after obtaining a new PMIxServer object

1 **Format**

PMIx v4.0 Python

```
2 rc = myserver.init(directives:list, map:dict)
```

Python

3 **IN directives**

4 List of Python **info** dictionaries (list)

5 **IN map**

6 Python dictionary key-function pairs that map **server module** callback functions to

7 provided implementations (dict)

8 Returns:

9 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

10 See **PMIx_server_init** for description of all relevant attributes and behaviors

11 A.4.2 Server.finalize

12 **Summary**

13 Finalize the PMIx server library

14 **Format**

PMIx v4.0 Python

```
15 rc = myserver.finalize()
```

Python

16 Returns:

17 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

18 See **PMIx_server_finalize** for details

19 A.4.3 Server.generate_regex

20 **Summary**

21 Generate a regular expression representation of the input strings.

1
PMIx v4.0

Format

Python

2 `rc, regex = myserver.generate_regex(input:list)`

Python

3 **IN** `input`

List of Python strings (e.g., node names) (list)

5 Returns:

- 6 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 7 • `regex` - Python **bytearray** containing regular expression representation of the input list
- 8 (**bytearray**)

9 See [PMIx_generate_regex](#) for details

10 A.4.4 Server.generate_ppn

11 Summary

12 Generate a regular expression representation of the input strings.

13 Format

PMIx v4.0

Python

14 `rc, regex = myserver.generate_ppn(input:list)`

Python

15 **IN** `input`

16 List of Python strings describing the ranks on each node (list)

17 Returns:

- 18 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 19 • `regex` - Python **bytearray** containing regular expression representation of the input list
- 20 (**bytearray**)

21 See [PMIx_generate_ppn](#) for details

22 A.4.5 Server.register_namespace

23 Summary

24 Setup the data about a particular namespace.

1
PMIx v4.0

Format

Python

```
2 rc = myserver.register_namespace(namespace:str,  
3                               nlocalprocs:integer,  
4                               directives:list)
```

Python

5 **IN namespace**
6 Python string containing the namespace (str)

7 **IN nlocalprocs**
8 Number of local processes (integer)

9 **IN directives**
10 List of Python [info](#) dictionaries (list)

11 Returns:

- 12 • *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

13 See [PMIx_server_register_namespace](#) for description of all relevant attributes and behaviors

14 A.4.6 Server.deregister_namespace

15 Summary

16 Deregister a namespace.

17 Format

PMIx v4.0

Python

```
18 myserver.deregister_namespace(namespace:str)
```

Python

19 **IN namespace**
20 Python string containing the namespace (str)

21 Returns: None

22 See [PMIx_server_deregister_namespace](#) for details

23 A.4.7 Server.register_client

24 Summary

25 Register a client process with the PMIx server library.

1 **Format**
PMIx v4.0 Python

2 `rc = myserver.register_client(proc:dict, uid:integer, gid:integer)`
Python

3 **IN** `proc`
4 Python `proc` dictionary identifying the client process (dict)

5 **IN** `uid`
6 Linux uid value for user executing client process (integer)

7 **IN** `gid`
8 Linux gid value for user executing client process (integer)

9 Returns:

10 • `rc - PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

11 See `PMIx_server_register_client` for details

12 **A.4.8 Server.deregister_client**

13 **Summary**
14 Deregister a client process and purge all data relating to it

15 **Format**
PMIx v4.0 Python

16 `myserver.deregister_client(proc:dict)`
Python

17 **IN** `proc`
18 Python `proc` dictionary identifying the client process (dict)

19 Returns: None

20 See `PMIx_server_deregister_client` for details

21 **A.4.9 Server.setup_fork**

22 **Summary**
23 Setup the environment of a child process that is to be forked by the host

1 **Format**
PMIx v4.0 Python

```
2 rc = myserver.setup_fork(proc:dict, environ:dict)
```

3 **IN** `proc`
4 Python `proc` dictionary identifying the client process (dict)

5 **INOUT** `environ`
6 Python dictionary containing the environment to be passed to the client (dict)

7 Returns:

8 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

9 See `PMIx_server_setup_fork` for details

10 A.4.10 Server.dmodex_request

11 **Summary**
12 Function by which the host server can request modex data from the local PMIx server.

13 **Format**
PMIx v4.0 Python

```
14 rc, data = myserver.dmodex_request(proc:dict)
```

15 **IN** `proc`
16 Python `proc` dictionary identifying the process whose data is requested (dict)

17 Returns:

18 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

19 • `data` - Python `byteobject` containing the returned data (dict)

20 See `PMIx_server_dmodex_request` for details

21 A.4.11 Server.setup_application

22 **Summary**
23 Function by which the resource manager can request application-specific setup data prior to launch
24 of a `job`.

1 **Format**

PMIx v4.0 Python

2 `rc, info = myserver.setup_application(namespace:str, directives:list)`
Python

3 **IN namespace**
4 Namespace whose setup information is being requested (str)

5 **IN directives**
6 Python list of **info** directives

7 Returns:

8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 • *info* - Python list of **info** dictionaries containing the returned data (list)

10 See **PMIx_server_setup_application** for details

11 A.4.12 Server.register_attributes

12 **Summary**
13 Register host environment attribute support for a function.

14 **Format**

PMIx v4.0 Python

15 `rc = myserver.register_attributes(function:str, attrs:list)`
Python

16 **IN function**
17 Name of the function (str)

18 **IN attrs**
19 Python list of **regattr** dictionaries describing the supported attributes

20 Returns:

21 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

22 See **PMIx_Register_attributes** for details

23 A.4.13 Server.setup_local_support

24 **Summary**
25 Function by which the local PMIx server can perform any application-specific operations prior to
26 spawning local clients of a given application

1 **Format**

PMIx v4.0 Python

```
2 rc = myserver.setup_local_support(namespace:str, info:list)
```

Python

3 **IN namespace**
 4 Namespace whose setup information is being requested (str)

5 **IN info**
 6 Python list of **info** dictionaries containing the setup data (list)

7 Returns:

8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 See **PMIx_server_setup_local_support** for details

10 A.4.14 Server.iof_deliver

11 **Summary**

12 Function by which the host environment can pass forwarded IO to the PMIx server library for
 13 distribution to its clients.

14 **Format**

PMIx v4.0 Python

```
15 rc = myserver.iof_deliver(source:dict, channel:integer,  
16                          data:dict, directives:list)
```

Python

17 **IN source**
 18 Python **proc** dictionary identifying the process who generated the data (dict)

19 **IN channel**
 20 Python **channel** bitmask identifying IO channel of the provided data (integer)

21 **IN data**
 22 Python **byteobject** containing the data (dict)

23 **IN directives**
 24 Python list of **info** dictionaries containing directives (list)

25 Returns:

26 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

27 See **PMIx_server_IOF_deliver** for details

28 A.4.15 Server.collect_inventory

29 **Summary**

30 Collect inventory of resources on a node

1 **Format**

PMIx v4.0 Python

```
2 rc, info = myserver.collect_inventory(directives:list)
```

Python

3 **IN directives**

4 Python list of **info** dictionaries containing directives (list)

5 Returns:

6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

7 • *info* - Python list of **info** dictionaries containing the returned data (list)

8 See **PMIx_server_collect_inventory** for details

9 A.4.16 Server.deliver_inventory

10 **Summary**

11 Pass collected inventory to the PMIx server library for storage

12 **Format**

PMIx v4.0 Python

```
13 rc = myserver.deliver_inventory(info:list, directives:list)
```

Python

14 **IN info**

15 - Python list of **info** dictionaries containing the inventory data (list)

16 **IN directives**

17 Python list of **info** dictionaries containing directives (list)

18 Returns:

19 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

20 See **PMIx_server_deliver_inventory** for details

21 A.5 PMIxTool

22 The tool Python class inherits the Python "server" class as its parent. Thus, it includes all client and
 23 server functions in addition to the ones defined in this section.

24 A.5.1 Tool.init

25 **Summary**

26 Initialize the PMIx tool library after obtaining a new PMIxTool object

1 **Format**
PMIx v4.0 Python

2 `rc,proc = mytool.init(info:list)`
Python

3 **IN info**
4 List of Python **info** dictionaries (list)

5 Returns:

6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

7 • *proc* - a Python **proc** dictionary (dict)

8 See **PMIx_tool_init** for description of all relevant attributes and behaviors

9 **A.5.2 Tool.finalize**

10 **Summary**
11 Finalize the PMIx tool library, closing the connection to the server.

12 **Format**
PMIx v4.0 Python

13 `rc = mytool.finalize()`
Python

14 Returns:

15 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

16 See **PMIx_tool_finalize** for description of all relevant attributes and behaviors

17 **A.5.3 Tool.connect_to_server**

18 **Summary**
19 Switch connection from the current PMIx server to another one, or initialize a connection to a
20 specified server.

1 **Format**
PMIx v4.0 Python

```
2 rc,proc = mytool.connect_to_server(info:list)
```

Python

3 **IN info**
 4 List of Python **info** dictionaries (list)

5 Returns:

6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

7 • *proc* - a Python **proc** dictionary (dict)

8 See **PMIx_tool_connect_to_server** for description of all relevant attributes and behaviors

9 A.5.4 Tool.iof_pull

10 **Summary**
 11 Register to receive output forwarded from a remote process.

12 **Format**
PMIx v4.0 Python

```
13 rc,id = mytool.iof_pull(sources:list, channel:integer, directives:list, cbfunc:func)
```

Python

14 **IN sources**
 15 List of Python **proc** dictionaries of processes whose IO is being requested (list)

16 **IN channel**
 17 Python **channel** bitmask identifying IO channels to be forwarded (integer)

18 **IN directives**
 19 List of Python **info** dictionaries describing request (list)

20 **IN cbfunc**
 21 Python **iofcbfunc** to receive IO payloads (func)

22 Returns:

23 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

24 • *id* - PMIx reference identifier for request (integer)

25 See **PMIx_IOF_pull** for description of all relevant attributes and behaviors

26 A.5.5 Tool.iof_deregister

27 **Summary**
 28 Deregister from output forwarded from a remote process.

1 **Format**
PMIx v4.0 Python

```
2 rc = mytool.iof_deregister(id:integer, directives:list)
```

Python

3 **IN id**
4 PMIx reference identifier returned by pull request (list)

5 **IN directives**
6 List of Python **info** dictionaries describing request (list)

7 Returns:

8 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

9 See **PMIx_IOF_deregister** for description of all relevant attributes and behaviors

10 A.5.6 Tool.iof_push

11 **Summary**
12 Push data collected locally (typically from stdin) to stdin of target recipients

13 **Format**
PMIx v4.0 Python

```
14 rc = mytool.iof_push(targets:list, data:dict, directives:list)
```

Python

15 **IN sources**
16 List of Python **proc** dictionaries of target processes (list)

17 **IN data**
18 Python **byteobject** dictionary containing data to be delivered (dict)

19 **IN directives**
20 List of Python **info** dictionaries describing request (list)

21 Returns:

22 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

23 See **PMIx_IOF_push** for description of all relevant attributes and behaviors

APPENDIX B

Acknowledgements

1 This document represents the work of many people who have contributed to the PMIx community.
2 Without the hard work and dedication of these people this document would not have been possible.
3 The sections below list some of the active participants and organizations in the various PMIx
4 standard iterations.

5 **B.1 Version 3.0**

6 The following list includes some of the active participants in the PMIx v3 standardization process.

- 7 • Ralph H. Castain, Andrew Friedley, Brandon Yates
- 8 • Joshua Hursey
- 9 • Aurelien Bouteiller and George Bosilca
- 10 • Dirk Schubert
- 11 • Kevin Harms

12 The following institutions supported this effort through time and travel support for the people listed
13 above.

- 14 • Intel Corporation
- 15 • IBM, Inc.
- 16 • University of Tennessee, Knoxville
- 17 • The Exascale Computing Project, an initiative of the US Department of Energy
- 18 • National Science Foundation
- 19 • Argonne National Laboratory
- 20 • Allinea (ARM)

1 **B.2 Version 2.0**

2 The following list includes some of the active participants in the PMIx v2 standardization process.

3 • Ralph H. Castain, Annapurna Dasari, Christopher A. Holguin, Andrew Friedley, Michael Klemm
4 and Terry Wilmarth

5 • Joshua Hursey, David Solt, Alexander Eichenberger, Geoff Paulsen, and Sameh Sharkawi

6 • Aurelien Bouteiller and George Bosilca

7 • Artem Polyakov, Igor Ivanov and Boris Karasev

8 • Gilles Gouaillardet

9 • Michael A Raymond and Jim Stoffel

10 • Dirk Schubert

11 • Moe Jette

12 • Takahiro Kawashima and Shinji Sumimoto

13 • Howard Pritchard

14 • David Beer

15 • Brice Goglin

16 • Geoffroy Vallee, Swen Boehm, Thomas Naughton and David Bernholdt

17 • Adam Moody and Martin Schulz

18 • Ryan Grant and Stephen Olivier

19 • Michael Karo

20 The following institutions supported this effort through time and travel support for the people listed
21 above.

22 • Intel Corporation

23 • IBM, Inc.

24 • University of Tennessee, Knoxville

25 • The Exascale Computing Project, an initiative of the US Department of Energy

26 • National Science Foundation

27 • Mellanox, Inc.

28 • Research Organization for Information Science and Technology

29 • HPE Co.

- 1 • Allinea (ARM)
- 2 • SchedMD, Inc.
- 3 • Fujitsu Limited
- 4 • Los Alamos National Laboratory
- 5 • Adaptive Solutions, Inc.
- 6 • INRIA
- 7 • Oak Ridge National Laboratory
- 8 • Lawrence Livermore National Laboratory
- 9 • Sandia National Laboratory
- 10 • Altair

11 **B.3 Version 1.0**

12 The following list includes some of the active participants in the PMIx v1 standardization process.

- 13 • Ralph H. Castain, Annapurna Dasari and Christopher A. Holguin
- 14 • Joshua Hursey and David Solt
- 15 • Aurelien Bouteiller and George Bosilca
- 16 • Artem Polyakov, Elena Shipunova, Igor Ivanov, and Joshua Ladd
- 17 • Gilles Gouaillardet
- 18 • Gary Brown
- 19 • Moe Jette

20 The following institutions supported this effort through time and travel support for the people listed
21 above.

- 22 • Intel Corporation
- 23 • IBM, Inc.
- 24 • University of Tennessee, Knoxville
- 25 • Mellanox, Inc.
- 26 • Research Organization for Information Science and Technology
- 27 • Adaptive Solutions, Inc.
- 28 • SchedMD, Inc.

Bibliography

- [1] Ralph H. Castain, David Solt, Joshua Hursey, and Aurelien Bouteiller. PMix: Process management for exascale environments. In *Proceedings of the 24th European MPI Users' Group Meeting, EuroMPI '17*, pages 14:1–14:10, New York, NY, USA, 2017. ACM.

Index

- application, [9](#), [12](#), [74](#), [75](#), [132](#), [190](#), [247](#), [249](#)
 - Defintion, [14](#)
- host environment
 - Defintion, [15](#)
- job, [9](#), [10](#), [12](#), [74–76](#), [132](#), [134](#), [190](#), [241](#),
[242](#), [246](#), [247](#), [249](#), [259](#), [261](#), [400](#)
 - Defintion, [14](#)
- namespace
 - Defintion, [14](#)
- network plane, [321](#)
 - Defintion, [15](#)
- network planes, [353](#)
 - Defintion, [15](#)
- PMIx_Abort, [8](#), [31](#), [154](#), [271](#), [272](#), [365](#), [377](#)
 - Defintion, [153](#)
- PMIX_ADD_ENVAR
 - Defintion, [87](#)
- PMIX_ADD_HOST, [156](#), [161](#), [285](#)
 - Defintion, [82](#)
- PMIX_ADD_HOSTFILE, [156](#), [161](#), [285](#)
 - Defintion, [82](#)
- PMIX_ALLOC_BANDWIDTH, [88](#), [193](#),
[196](#), [260](#), [305](#)
 - Defintion, [88](#)
- PMIX_ALLOC_CPU_LIST, [193](#), [196](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_DIRECTIVE, [68](#)
- PMIx_Alloc_directive_string, [9](#), [395](#)
 - Defintion, [110](#)
- pmix_alloc_directive_t, [47](#), [68](#), [110](#), [303](#),
[360](#), [394](#)
 - Defintion, [47](#)
- PMIX_ALLOC_EXTEND, [47](#)
- PMIX_ALLOC_EXTERNAL, [47](#)
- PMIX_ALLOC_ID, [192](#), [195](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_MEM_SIZE, [193](#), [196](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NETWORK, [193](#), [196](#), [260](#),
[304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NETWORK_ENDPTS, [88](#),
[193](#), [194](#), [196](#), [197](#), [260](#), [304](#)
 - Defintion, [89](#)
- PMIX_ALLOC_NETWORK_ENDPTS_NODE,
[194](#), [197](#), [260](#)
 - Defintion, [89](#)
- PMIX_ALLOC_NETWORK_ID, [88](#), [193](#),
[196](#), [260](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NETWORK_PLANE, [88](#),
[193](#), [194](#), [196](#), [197](#), [260](#), [304](#)
 - Defintion, [89](#)
- PMIX_ALLOC_NETWORK_QOS, [88](#), [193](#),
[194](#), [196](#), [260](#), [305](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NETWORK_SEC_KEY,
[88](#), [193](#), [194](#), [196](#), [197](#), [260](#), [305](#)
 - Defintion, [89](#)
- PMIX_ALLOC_NETWORK_TYPE, [88](#),
[193](#), [194](#), [196](#), [197](#), [260](#), [304](#)
 - Defintion, [89](#)
- PMIX_ALLOC_NEW, [47](#)
- PMIX_ALLOC_NODE_LIST, [193](#), [196](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NUM_CPU_LIST, [193](#),
[196](#), [304](#)
 - Defintion, [88](#)
- PMIX_ALLOC_NUM_CPUS, [193](#), [196](#), [304](#)

Defintion, [88](#)
 PMIX_ALLOC_NUM_NODES, [193](#), [195](#),
 [304](#)
 Defintion, [88](#)
 PMIX_ALLOC_REAQUIRE, [47](#)
 PMIX_ALLOC_RELEASE, [47](#)
 PMIX_ALLOC_TIME, [193](#), [196](#), [260](#), [304](#)
 Defintion, [88](#)
 PMIX_ALLOCATED_NODELIST, [243](#)
 Defintion, [74](#)
 PMIx_Allocation_request, [11](#), [191](#), [197](#)
 Defintion, [192](#)
 PMIx_Allocation_request_nb, [9](#), [88](#), [179](#),
 [197](#), [385](#)
 Defintion, [194](#)
 PMIX_ANL_MAP, [244](#)
 Defintion, [80](#)
 PMIX_APP, [67](#)
 PMIX_APP_CONSTRUCT
 Defintion, [53](#)
 PMIX_APP_CREATE
 Defintion, [53](#)
 PMIX_APP_DESTRUCT
 Defintion, [53](#)
 PMIX_APP_FREE
 Defintion, [53](#)
 PMIX_APP_INFO, [127](#), [130](#), [134](#), [182](#), [187](#)
 Defintion, [75](#)
 PMIX_APP_INFO_ARRAY, [76](#), [250](#)
 Defintion, [75](#)
 PMIX_APP_INFO_CREATE, [10](#), [11](#)
 Defintion, [54](#)
 PMIX_APP_MAP_REGEX
 Defintion, [80](#)
 PMIX_APP_MAP_TYPE
 Defintion, [80](#)
 PMIX_APP_RANK, [243](#)
 Defintion, [73](#)
 PMIX_APP_SIZE, [134](#), [243](#), [249](#)
 Defintion, [76](#)
 pmix_app_t, [10](#), [11](#), [52–54](#), [63](#), [66](#), [155](#), [159](#),
 [284](#), [361](#)
 Defintion, [52](#)
 PMIX_APPEND_ENVAR
 Defintion, [87](#)
 PMIX_APPLDR, [243](#), [249](#)
 Defintion, [74](#)
 PMIX_APPNUM, [75](#), [127](#), [130](#), [134](#), [182](#),
 [187](#), [242](#), [243](#), [250](#)
 Defintion, [73](#)
 PMIX_ARCH
 Defintion, [73](#)
 PMIX_ARGV_APPEND
 Defintion, [62](#)
 PMIX_ARGV_APPEND_UNIQUE
 Defintion, [63](#)
 PMIX_ARGV_COPY
 Defintion, [65](#)
 PMIX_ARGV_COUNT
 Defintion, [65](#)
 PMIX_ARGV_FREE
 Defintion, [64](#)
 PMIX_ARGV_JOIN
 Defintion, [65](#)
 PMIX_ARGV_SPLIT
 Defintion, [64](#)
 PMIX_ATTR_UNDEF
 Defintion, [68](#)
 PMIX_AVAIL_PHYS_MEMORY, [244](#)
 Defintion, [77](#)
 PMIX_BINDTO, [157](#), [161](#), [244](#), [286](#)
 Defintion, [82](#)
 PMIX_BOOL, [67](#)
 PMIX_BUFFER, [67](#)
 PMIX_BYTE, [67](#)
 PMIX_BYTE_OBJECT, [67](#)
 PMIX_BYTE_OBJECT_CREATE
 Defintion, [60](#)
 PMIX_BYTE_OBJECT_DESTRUCT
 Defintion, [59](#)
 PMIX_BYTE_OBJECT_FREE
 Defintion, [60](#)
 PMIX_BYTE_OBJECT_LOAD
 Defintion, [60](#)
 pmix_byte_object_t, [59](#), [60](#), [67](#), [105](#), [176](#),
 [236](#), [264](#), [313](#), [317](#), [359](#)

Defintion, [59](#)
 pmix_check_key
 Defintion, [26](#)
 pmix_check_nspace
 Defintion, [27](#)
 pmix_check_procid
 Defintion, [30](#)
 PMIX_CLEANUP_EMPTY, [199](#), [202](#)
 Defintion, [90](#)
 PMIX_CLEANUP_IGNORE, [199](#), [202](#)
 Defintion, [90](#)
 PMIX_CLEANUP_LEAVE_TOPDIR, [199](#),
 [202](#)
 Defintion, [90](#)
 PMIX_CLEANUP_RECURSIVE, [199](#), [202](#)
 Defintion, [90](#)
 PMIX_CLIENT_ATTRIBUTES, [12](#), [182](#),
 [187](#)
 Defintion, [92](#)
 PMIX_CLIENT_AVG_MEMORY
 Defintion, [77](#)
 PMIX_CLIENT_FUNCTIONS
 Defintion, [92](#)
 PMIX_CLUSTER_ID
 Defintion, [73](#)
 PMIX_COLLECT_DATA, [138](#), [140](#), [274](#)
 Defintion, [78](#)
 PMIX_COLLECTIVE_ALGO, [9](#), [138](#), [141](#),
 [165](#), [168](#), [274](#), [289](#)
 Defintion, [78](#)
 PMIX_COLLECTIVE_ALGO_REQD, [138](#),
 [141](#), [165](#), [168](#), [274](#), [289](#)
 Defintion, [79](#)
 PMIX_COMMAND, [68](#)
 PMIx_Commit, [8](#), [103](#), [124](#), [125](#), [137](#), [258](#),
 [277](#), [379](#)
 Defintion, [137](#)
 PMIX_COMPRESSED_STRING, [68](#)
 PMIx_Connect, [8](#), [9](#), [23](#), [158](#), [166](#), [168](#), [170](#),
 [325–327](#), [382](#)
 Defintion, [164](#)
 PMIX_CONNECT_MAX_RETRIES, [117](#)
 Defintion, [70](#)
 PMIx_Connect_nb, [8](#), [166](#)
 Defintion, [166](#)
 PMIX_CONNECT_REQUESTED, [23](#)
 PMIX_CONNECT_RETRY_DELAY, [116](#)
 Defintion, [70](#)
 PMIX_CONNECT_SYSTEM_FIRST, [116](#),
 [118](#), [120](#)
 Defintion, [70](#)
 PMIX_CONNECT_TO_SYSTEM, [116](#), [118](#),
 [120](#)
 Defintion, [69](#)
 pmix_connection_cbfunc_t, [296](#)
 Defintion, [103](#)
 PMIX_COORD, [68](#)
 PMIX_COORD_CONSTRUCT
 Defintion, [353](#)
 PMIX_COORD_CREATE
 Defintion, [354](#)
 PMIX_COORD_DESTRUCT
 Defintion, [353](#)
 PMIX_COORD_FREE
 Defintion, [354](#)
 PMIX_COORD_LOGICAL_VIEW, [354](#)
 PMIX_COORD_PHYSICAL_VIEW, [354](#)
 pmix_coord_t, [68](#), [352–354](#)
 Defintion, [352](#)
 pmix_coord_view_t, [355](#)
 Defintion, [354](#)
 PMIX_COORD_VIEW_UNDEF, [354](#)
 PMIX_COSPAWN_APP
 Defintion, [83](#)
 PMIX_CPU_LIST, [158](#), [162](#), [287](#)
 Defintion, [84](#)
 PMIX_CPUS_PER_PROC, [157](#), [162](#), [287](#)
 Defintion, [83](#)
 PMIX_CPuset
 Defintion, [72](#)
 PMIX_CRED_TYPE, [312](#)
 Defintion, [91](#)
 PMIX_CREDENTIAL
 Defintion, [72](#)
 pmix_credential_cbfunc_t, [234](#), [311](#)
 Defintion, [104](#)

PMIX_CRYPTO_KEY
 Defintion, [91](#)
 PMIX_DAEMON_MEMORY
 Defintion, [77](#)
 PMIX_DATA_ARRAY, [68](#)
 PMIX_DATA_ARRAY_CONSTRUCT
 Defintion, [35](#), [61](#)
 PMIX_DATA_ARRAY_CREATE
 Defintion, [36](#), [62](#)
 PMIX_DATA_ARRAY_DESTRUCT
 Defintion, [35](#), [61](#)
 PMIX_DATA_ARRAY_FREE
 Defintion, [36](#)
 PMIX_DATA_ARRAY_RELEASE
 Defintion, [62](#)
 pmix_data_array_t, [10](#), [11](#), [35](#), [36](#), [61](#), [62](#),
 [68](#), [84](#), [88](#), [184](#), [189](#), [191](#), [193](#),
 [196](#), [248–251](#), [260](#), [298](#), [304](#), [321](#),
 [353](#), [355](#), [356](#), [360](#)
 Defintion, [35](#), [61](#)
 PMIX_DATA_BUFFER_CONSTRUCT,
 [227](#), [229](#)
 Defintion, [224](#)
 PMIX_DATA_BUFFER_CREATE, [227](#),
 [229](#)
 Defintion, [224](#)
 PMIX_DATA_BUFFER_DESTRUCT
 Defintion, [225](#)
 PMIX_DATA_BUFFER_LOAD
 Defintion, [225](#)
 PMIX_DATA_BUFFER_RELEASE
 Defintion, [224](#)
 pmix_data_buffer_t, [223–228](#), [232](#)
 Defintion, [223](#)
 PMIX_DATA_BUFFER_UNLOAD, [239](#),
 [240](#)
 Defintion, [226](#)
 PMIx_Data_copy, [9](#)
 Defintion, [230](#)
 PMIx_Data_copy_payload, [9](#)
 Defintion, [231](#)
 PMIx_Data_pack, [9](#), [227](#), [239](#), [240](#)
 Defintion, [226](#)
 PMIx_Data_print, [9](#)
 Defintion, [230](#)
 PMIX_DATA_RANGE, [68](#)
 PMIx_Data_range_string, [9](#), [393](#)
 Defintion, [109](#)
 pmix_data_range_t, [34](#), [68](#), [109](#), [221](#), [295](#),
 [359](#), [393](#)
 Defintion, [34](#)
 PMIX_DATA_SCOPE, [126](#), [130](#)
 Defintion, [79](#)
 PMIX_DATA_TYPE, [68](#)
 PMIx_Data_type_string, [9](#), [394](#)
 Defintion, [109](#)
 pmix_data_type_t, [35](#), [36](#), [39](#), [41](#), [43](#), [51](#), [58](#),
 [61](#), [62](#), [67](#), [68](#), [109](#), [227](#), [229–231](#),
 [358](#), [394](#)
 Defintion, [67](#)
 PMIx_Data_unpack, [9](#)
 Defintion, [228](#)
 PMIX_DEBUG_APP_DIRECTIVES
 Defintion, [87](#)
 PMIX_DEBUG_JOB
 Defintion, [87](#)
 PMIX_DEBUG_JOB_DIRECTIVES
 Defintion, [87](#)
 PMIX_DEBUG_STOP_IN_INIT
 Defintion, [86](#)
 PMIX_DEBUG_STOP_ON_EXEC
 Defintion, [86](#)
 PMIX_DEBUG_WAIT_FOR_NOTIFY
 Defintion, [87](#)
 PMIX_DEBUG_WAITING_FOR_NOTIFY
 Defintion, [87](#)
 PMIX_DEBUGGER_DAEMONS, [157](#), [162](#),
 [286](#)
 Defintion, [83](#)
 PMIx_Deregister_event_handler, [9](#), [391](#)
 Defintion, [219](#)
 PMIx_Disconnect, [8](#), [9](#), [23](#), [166](#), [170](#), [172](#),
 [327](#), [383](#)
 Defintion, [168](#)
 PMIx_Disconnect_nb, [8](#), [172](#), [327](#)
 Defintion, [170](#)

PMIX_DISPLAY_MAP, 156, 161, 286
 Defintion, 82
 pmix_dmodex_response_fn_t, 258
 Defintion, 102
 PMIX_DOUBLE, 67
 PMIX_DSTPATH
 Defintion, 70
 PMIX_EMBED_BARRIER, 115
 Defintion, 79
 PMIX_ENUM_VALUE, 12, 56
 Defintion, 92
 PMIX_ENVAR, 68
 PMIX_ENVAR_CONSTRUCT
 Defintion, 48
 PMIX_ENVAR_CREATE
 Defintion, 48
 PMIX_ENVAR_DESTRUCT
 Defintion, 48
 PMIX_ENVAR_FREE
 Defintion, 49
 PMIX_ENVAR_LOAD
 Defintion, 49
 pmix_envar_t, 48, 49, 68, 360
 Defintion, 47
 PMIX_ERR_BAD_PARAM, 22
 PMIX_ERR_COMM_FAILURE, 22
 PMIX_ERR_DATA_VALUE_NOT_FOUND,
 22
 PMIX_ERR_DEBUGGER_RELEASE, 21
 PMIX_ERR_DUPLICATE_KEY, 23
 PMIX_ERR_EVENT_REGISTRATION, 23
 PMIX_ERR_HANDSHAKE_FAILED, 21
 PMIX_ERR_IN_ERRNO, 22
 PMIX_ERR_INIT, 22
 PMIX_ERR_INVALID_ARG, 22
 PMIX_ERR_INVALID_ARGS, 22
 PMIX_ERR_INVALID_CRED, 21
 PMIX_ERR_INVALID_KEY, 22
 PMIX_ERR_INVALID_KEY_LENGTH, 22
 PMIX_ERR_INVALID_KEYVALP, 22
 PMIX_ERR_INVALID_LENGTH, 22
 PMIX_ERR_INVALID_NAMESPACE, 22
 PMIX_ERR_INVALID_NUM_ARGS, 22
 PMIX_ERR_INVALID_NUM_PARSED, 22
 PMIX_ERR_INVALID_OPERATION, 23
 PMIX_ERR_INVALID_SIZE, 22
 PMIX_ERR_INVALID_TERMINATION,
 23
 PMIX_ERR_INVALID_VAL, 22
 PMIX_ERR_INVALID_VAL_LENGTH, 22
 PMIX_ERR_IOF_COMPLETE, 24
 PMIX_ERR_IOF_FAILURE, 24
 PMIX_ERR_JOB_TERMINATED, 23
 PMIX_ERR_LOST_CONNECTION_TO_CLIENT,
 22
 PMIX_ERR_LOST_CONNECTION_TO_SERVER,
 22
 PMIX_ERR_LOST_PEER_CONNECTION,
 22
 PMIX_ERR_NO_PERMISSIONS, 22
 PMIX_ERR_NODE_DOWN, 24
 PMIX_ERR_NODE_OFFLINE, 24
 PMIX_ERR_NOMEM, 22
 PMIX_ERR_NOT_FOUND, 22
 PMIX_ERR_NOT_IMPLEMENTED, 22
 PMIX_ERR_NOT_SUPPORTED, 22
 PMIX_ERR_OUT_OF_RESOURCE, 22
 PMIX_ERR_PACK_FAILURE, 22
 PMIX_ERR_PACK_MISMATCH, 22
 PMIX_ERR_PARTIAL_SUCCESS, 23
 PMIX_ERR_PROC_ABORTED, 21
 PMIX_ERR_PROC_ABORTING, 21
 PMIX_ERR_PROC_CHECKPOINT, 21
 PMIX_ERR_PROC_ENTRY_NOT_FOUND,
 21
 PMIX_ERR_PROC_MIGRATE, 21
 PMIX_ERR_PROC_REQUESTED_ABORT,
 21
 PMIX_ERR_PROC_RESTART, 21
 PMIX_ERR_READY_FOR_HANDSHAKE,
 21
 PMIX_ERR_REPEAT_ATTR_REGISTRATION,
 24
 PMIX_ERR_RESOURCE_BUSY, 22
 PMIX_ERR_SERVER_FAILED_REQUEST,
 21

PMIX_ERR_SERVER_NOT_AVAIL, 22
 PMIX_ERR_SILENT, 21
 PMIX_ERR_SYS_BASE, 24
 PMIX_ERR_SYS_OTHER, 25
 PMIX_ERR_TIMEOUT, 22
 PMIX_ERR_TYPE_MISMATCH, 21
 PMIX_ERR_UNKNOWN_DATA_TYPE, 21
 PMIX_ERR_UNPACK_FAILURE, 22
 PMIX_ERR_UNPACK_INADEQUATE_SPACE, 21
 PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER, 22
 PMIX_ERR_UNREACH, 22
 PMIX_ERR_UPDATE_ENDPOINTS, 23
 PMIX_ERR_WOULD_BLOCK, 21
 PMIX_ERROR, 21
 PMIx_Error_string, 8, 392
 Defintion, 108
 PMIX_EVENT_ACTION_COMPLETE, 25
 PMIX_EVENT_ACTION_DEFERRED, 25
 PMIX_EVENT_ACTION_TIMEOUT, 218
 Defintion, 81
 PMIX_EVENT_AFFECTED_PROC, 218, 222
 Defintion, 81
 PMIX_EVENT_AFFECTED_PROCS, 218, 222
 Defintion, 81
 PMIX_EVENT_BASE, 114, 117, 123
 Defintion, 69
 PMIX_EVENT_CUSTOM_RANGE, 218, 222
 Defintion, 81
 PMIX_EVENT_DO_NOT_CACHE
 Defintion, 81
 PMIX_EVENT_HDLR_AFTER, 217
 Defintion, 80
 PMIX_EVENT_HDLR_APPEND, 218
 Defintion, 81
 PMIX_EVENT_HDLR_BEFORE, 217
 Defintion, 80
 PMIX_EVENT_HDLR_FIRST, 217
 Defintion, 80
 PMIX_EVENT_HDLR_FIRST_IN_CATEGORY, 217
 Defintion, 80
 PMIX_EVENT_HDLR_LAST, 217
 Defintion, 80
 PMIX_EVENT_HDLR_LAST_IN_CATEGORY, 217
 Defintion, 80
 PMIX_EVENT_HDLR_NAME, 217
 Defintion, 80
 PMIX_EVENT_HDLR_PREPEND, 217
 Defintion, 80
 PMIX_EVENT_NO_ACTION_TAKEN, 25
 PMIX_EVENT_NO_TERMINATION
 Defintion, 81
 PMIX_EVENT_NON_DEFAULT, 222
 Defintion, 81
 pmix_event_notification_cbfunc_fn_t, 99, 100
 Defintion, 99
 PMIX_EVENT_PARTIAL_ACTION_TAKEN, 25
 PMIX_EVENT_PROXY
 Defintion, 81
 PMIX_EVENT_RETURN_OBJECT, 218
 Defintion, 81
 PMIX_EVENT_SILENT_TERMINATION, 218
 Defintion, 81
 PMIX_EVENT_TERMINATE_JOB, 218
 Defintion, 81
 PMIX_EVENT_TERMINATE_NODE, 218
 Defintion, 81
 PMIX_EVENT_TERMINATE_PROC, 218
 Defintion, 81
 PMIX_EVENT_TERMINATE_SESSION, 218
 Defintion, 81
 PMIX_EVENT_TEXT_MESSAGE
 Defintion, 81
 PMIX_EVENT_WANT_TERMINATION
 Defintion, 81

pmix_evhdlr_reg_cbfunc_t, 98, 217
 Defintion, **98**
 PMIX_EXISTS, 21
 PMIX_EXIT_CODE
 Defintion, **74**
 PMIX_EXTERNAL_ERR_BASE, 25
 pmix_fabric_t, 319–323
 Defintion, **319**
 PMIX_FABRIC_UPDATE_PENDING, 320
 PMIX_FABRIC_UPDATED, 320
 PMIx_Fence, 3, 6, 8, 13, 123, 124, 139, 141,
 166, 170, 258, 272, 275, 325, 331,
 336, 365, 379
 Defintion, **137**
 PMIx_Fence_nb, 8, 12, 96, 141, 272, 275,
 365
 Defintion, **139**
 PMIx_Finalize, 8, 23, 31, 79, 114, 115, 163,
 270, 271, 364, 377
 Defintion, **115**
 PMIX_FLOAT, 67
 PMIX_FWD_ALL_CHANNELS, 47
 PMIX_FWD_NO_CHANNELS, 47
 PMIX_FWD_STDDIAG, 11
 Defintion, **83**
 PMIX_FWD_STDDIAG_CHANNEL, 47
 PMIX_FWD_STDERR, 157, 162, 286, 300
 Defintion, **83**
 PMIX_FWD_STDERR_CHANNEL, 47
 PMIX_FWD_STDIN, 157, 162, 286, 300
 Defintion, **83**
 PMIX_FWD_STDIN_CHANNEL, 47
 PMIX_FWD_STDOUT, 157, 162, 286, 300
 Defintion, **83**
 PMIX_FWD_STDOUT_CHANNEL, 47
 PMIX_GDS_ACTION_COMPLETE, 23
 PMIX_GDS_MODULE, 114, 117, 123
 Defintion, **72**
 PMIx_generate_ppn, 8, 397
 Defintion, **239**
 PMIx_generate_regex, 8, 240, 246, 397
 Defintion, **238**
 PMIx_Get, 3, 8–10, 36, 68–72, 76, 78–90,
 114, 126, 128, 130–134, 136,
 156–158, 160–163, 185, 190, 191,
 241, 243, 244, 267, 285–287, 324,
 331, 353, 380
 Defintion, **125**
 PMIx_Get_credential, 11, 91, 312, 387
 Defintion, **234**
 PMIx_Get_nb, 8, 17, 97
 Defintion, **128**
 PMIx_Get_version, 8, 17, 376
 Defintion, **112**
 PMIX_GLOBAL, 33
 PMIX_GLOBAL_RANK, 244
 Defintion, **73**
 PMIX_GROUP_ACCEPT, 59
 PMIX_GROUP_ASSIGN_CONTEXT_ID,
 329, 333, 340, 343
 Defintion, **93**
 PMIx_Group_construct, 326, 330, 331, 334,
 388
 Defintion, **327**
 PMIX_GROUP_CONSTRUCT_ABORT, 24
 PMIX_GROUP_CONSTRUCT_COMPLETE,
 24
 PMIx_Group_construct_nb, 334
 Defintion, **331**
 PMIX_GROUP_CONTEXT_ID
 Defintion, **93**
 PMIX_GROUP_CONTEXT_ID_ASSIGNED,
 24
 PMIX_GROUP_DECLINE, 59
 PMIx_Group_destruct, 327, 336, 338, 390
 Defintion, **334**
 PMIx_Group_destruct_nb, 338
 Defintion, **336**
 PMIX_GROUP_ENDPT_DATA
 Defintion, **93**
 PMIX_GROUP_ID
 Defintion, **93**
 PMIx_Group_invite, 326, 341, 342, 344, 388
 Defintion, **338**
 PMIX_GROUP_INVITE_ACCEPTED, 24
 PMIX_GROUP_INVITE_DECLINE

Defintion, [93](#)
 PMIX_GROUP_INVITE_DECLINED, [24](#)
 PMIX_GROUP_INVITE_FAILED, [24](#)
 PMIx_Group_invite_nb
 Defintion, [342](#)
 PMIX_GROUP_INVITED, [23](#)
 PMIx_Group_join, [59](#), [326](#), [341](#), [344](#), [346](#),
 [347](#), [349](#), [389](#)
 Defintion, [344](#)
 PMIx_Group_join_nb, [344](#), [349](#)
 Defintion, [347](#)
 PMIX_GROUP_LEADER, [328](#), [330](#), [333](#),
 [342](#), [347](#)
 Defintion, [93](#)
 PMIX_GROUP_LEADER_FAILED, [24](#)
 PMIX_GROUP_LEADER_SELECTED, [24](#)
 PMIx_Group_leave, [327](#), [350](#), [351](#), [389](#)
 Defintion, [349](#)
 PMIx_Group_leave_nb
 Defintion, [350](#)
 PMIX_GROUP_LEFT, [24](#)
 PMIX_GROUP_LOCAL_ONLY, [329](#), [333](#)
 Defintion, [93](#)
 PMIX_GROUP_MEMBER_FAILED, [24](#)
 PMIX_GROUP_MEMBERSHIP, [330](#)
 Defintion, [93](#)
 PMIX_GROUP_MEMBERSHIP_UPDATE,
 [24](#)
 PMIX_GROUP_NOTIFY_TERMINATION,
 [329](#), [330](#), [333](#), [336](#), [340](#), [343](#)
 Defintion, [93](#)
 pmix_group_opt_t, [59](#), [345](#), [348](#), [389](#)
 Defintion, [59](#)
 PMIX_GROUP_OPTIONAL, [328](#), [330](#), [333](#),
 [339](#), [343](#)
 Defintion, [93](#)
 PMIX_GRPID, [106](#), [142](#), [144](#), [146](#), [148](#),
 [150](#), [152](#), [183](#), [188](#), [192](#), [195](#), [198](#),
 [201](#), [204](#), [206](#), [209](#), [211](#), [235](#), [237](#),
 [278](#)–[283](#), [285](#), [292](#), [298](#), [300](#), [301](#),
 [303](#), [306](#), [309](#), [311](#), [313](#), [315](#), [318](#)
 Defintion, [70](#)
 pmix_hdlr_reg_cbfunc_t, [173](#), [175](#)

Defintion, [107](#)
 PMIx_Heartbeat, [9](#)
 Defintion, [207](#)
 PMIX_HOST, [156](#), [160](#), [285](#)
 Defintion, [82](#)
 PMIX_HOST_ATTRIBUTES, [12](#), [183](#), [187](#),
 [191](#)
 Defintion, [92](#)
 PMIX_HOST_FUNCTIONS
 Defintion, [92](#)
 PMIX_HOSTFILE, [156](#), [160](#), [285](#)
 Defintion, [82](#)
 PMIX_HOSTNAME, [75](#), [76](#), [127](#), [130](#), [136](#),
 [182](#), [187](#), [244](#)
 Defintion, [74](#)
 PMIX_HWLOC_HOLE_KIND
 Defintion, [78](#)
 PMIX_HWLOC_SHARE_TOPO
 Defintion, [78](#)
 PMIX_HWLOC_SHMEM_ADDR
 Defintion, [78](#)
 PMIX_HWLOC_SHMEM_FILE
 Defintion, [78](#)
 PMIX_HWLOC_SHMEM_SIZE
 Defintion, [78](#)
 PMIX_HWLOC_XML_V1, [244](#)
 Defintion, [78](#)
 PMIX_HWLOC_XML_V2, [244](#)
 Defintion, [78](#)
 PMIX_IMMEDIATE, [126](#), [130](#)
 Defintion, [78](#)
 PMIX_INDEX_ARGV, [157](#), [162](#), [287](#)
 Defintion, [83](#)
 PMIX_INFO, [67](#)
 PMIX_INFO_ARRAY_END, [45](#)
 pmix_info_cbfunc_t, [94](#), [97](#), [186](#), [195](#), [198](#),
 [200](#), [201](#), [203](#), [206](#), [207](#), [265](#), [297](#),
 [303](#), [305](#), [308](#), [332](#), [342](#), [348](#)
 Defintion, [97](#)
 PMIX_INFO_CONSTRUCT
 Defintion, [41](#)
 PMIX_INFO_CREATE, [45](#), [46](#)
 Defintion, [42](#)

PMIX_INFO_DESTRUCT
 Defintion, [42](#)
 PMIX_INFO_DIRECTIVES, [68](#)
 PMIx_Info_directives_string, [9](#), [394](#)
 Defintion, [109](#)
 pmix_info_directives_t, [44](#), [45](#), [68](#), [109](#), [360](#),
 [394](#)
 Defintion, [44](#)
 PMIX_INFO_FREE
 Defintion, [42](#)
 PMIX_INFO_IS_END, [10](#), [12](#)
 Defintion, [46](#)
 PMIX_INFO_IS_OPTIONAL
 Defintion, [46](#)
 PMIX_INFO_IS_REQUIRED, [44](#), [45](#)
 Defintion, [46](#)
 PMIX_INFO_LOAD
 Defintion, [43](#)
 PMIX_INFO_OPTIONAL
 Defintion, [45](#)
 PMIX_INFO_REQD, [45](#)
 PMIX_INFO_REQUIRED, [44](#)
 Defintion, [45](#)
 pmix_info_t, [3](#), [9](#), [10](#), [12](#), [13](#), [26](#), [34](#), [41–47](#),
 [54–56](#), [58](#), [68](#), [75](#), [76](#), [86](#), [88](#), [89](#),
 [98–100](#), [106](#), [107](#), [113–116](#), [121](#),
 [123](#), [143](#), [147](#), [173](#), [175](#), [176](#), [181](#),
 [185](#), [191–193](#), [195](#), [196](#), [200](#), [203](#),
 [207](#), [210](#), [213](#), [221](#), [234](#), [236](#), [246](#),
 [248–251](#), [260](#), [264–266](#), [295](#), [299](#),
 [300](#), [302](#), [304](#), [305](#), [308](#), [315](#), [328](#),
 [330](#), [332](#), [335](#), [337](#), [339](#), [342](#), [345](#),
 [348–350](#), [360](#), [362](#)
 Defintion, [41](#)
 PMIX_INFO_TRUE
 Defintion, [44](#)
 PMIX_INFO_XFER, [246](#)
 Defintion, [43](#)
 PMIx_Init, [9](#), [83](#), [86](#), [87](#), [112](#), [114](#), [115](#), [157](#),
 [161](#), [269](#), [286](#), [376](#)
 Defintion, [112](#)
 PMIx_Initialized, [8](#), [376](#)
 Defintion, [111](#)
 PMIX_INT, [67](#)
 PMIX_INT16, [67](#)
 PMIX_INT32, [67](#)
 PMIX_INT64, [67](#)
 PMIX_INT8, [67](#)
 PMIX_INTERNAL, [33](#)
 PMIX_IOF_BUFFERING_SIZE, [174](#), [177](#),
 [316](#)
 Defintion, [91](#)
 PMIX_IOF_BUFFERING_TIME, [174](#), [177](#),
 [316](#)
 Defintion, [91](#)
 PMIX_IOF_CACHE_SIZE, [174](#), [177](#), [316](#)
 Defintion, [91](#)
 pmix_iof_cbfunc_t, [173](#), [363](#)
 Defintion, [106](#)
 PMIX_IOF_CHANNEL, [68](#)
 PMIx_IOF_channel_string, [11](#), [395](#)
 Defintion, [110](#)
 pmix_iof_channel_t, [47](#), [68](#), [107](#), [110](#), [173](#),
 [264](#), [315](#), [360](#), [395](#)
 Defintion, [47](#)
 PMIX_IOF_COMPLETE, [107](#), [363](#)
 Defintion, [91](#)
 PMIx_IOF_deregister, [11](#), [406](#)
 Defintion, [175](#)
 PMIX_IOF_DROP_NEWEST, [174](#), [177](#), [316](#)
 Defintion, [91](#)
 PMIX_IOF_DROP_OLDEST, [174](#), [177](#), [316](#)
 Defintion, [91](#)
 PMIx_IOF_pull, [11](#), [175](#), [405](#)
 Defintion, [173](#)
 PMIx_IOF_push, [11](#), [406](#)
 Defintion, [176](#)
 PMIX_IOF_TAG_OUTPUT, [174](#)
 Defintion, [91](#)
 PMIX_IOF_TIMESTAMP_OUTPUT, [174](#)
 Defintion, [92](#)
 PMIX_IOF_XML_OUTPUT, [174](#)
 Defintion, [92](#)
 PMIX_JCTRL_CHECKPOINT, [22](#)
 PMIX_JCTRL_CHECKPOINT_COMPLETE,
 [22](#)

PMIX_JCTRL_PREEMPT_ALERT, 22
 PMIX_JOB_CONTINUOUS, 158, 162, 287
 Defintion, **84**
 PMIx_Job_control, 10, 89, 199, 202, 203,
 307
 Defintion, **198**
 PMIx_Job_control_nb, 9, 89, 179, 197, 245,
 386
 Defintion, **200**
 PMIX_JOB_CTRL_CANCEL, 199, 202,
 307
 Defintion, **89**
 PMIX_JOB_CTRL_CHECKPOINT, 199,
 203, 307
 Defintion, **89**
 PMIX_JOB_CTRL_CHECKPOINT_EVENT,
 199, 203, 307
 Defintion, **89**
 PMIX_JOB_CTRL_CHECKPOINT_METHOD,
 200, 203, 308
 Defintion, **89**
 PMIX_JOB_CTRL_CHECKPOINT_SIGNAL,
 200, 203, 307
 Defintion, **89**
 PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT,
 200, 203, 307
 Defintion, **89**
 PMIX_JOB_CTRL_ID, 89, 198, 199, 201,
 202, 307
 Defintion, **89**
 PMIX_JOB_CTRL_KILL, 199, 202, 307
 Defintion, **89**
 PMIX_JOB_CTRL_PAUSE, 199, 202, 307
 Defintion, **89**
 PMIX_JOB_CTRL_PREEMPTIBLE, 200,
 203, 308
 Defintion, **90**
 PMIX_JOB_CTRL_PROVISION, 200, 203,
 308
 Defintion, **90**
 PMIX_JOB_CTRL_PROVISION_IMAGE,
 200, 203, 308
 Defintion, **90**
 PMIX_JOB_CTRL_RESTART, 199, 202,
 307
 Defintion, **89**
 PMIX_JOB_CTRL_RESUME, 199, 202,
 307
 Defintion, **89**
 PMIX_JOB_CTRL_SIGNAL, 199, 202, 307
 Defintion, **89**
 PMIX_JOB_CTRL_TERMINATE, 199,
 202, 307
 Defintion, **90**
 PMIX_JOB_INFO, 126, 130, 134, 182, 187
 Defintion, **75**
 PMIX_JOB_INFO_ARRAY, 10, 76, 249
 Defintion, **75**
 PMIX_JOB_NUM_APPS, 134, 243, 249
 Defintion, **76**
 PMIX_JOB_RECOVERABLE, 158, 162,
 287
 Defintion, **84**
 PMIX_JOB_SIZE, 10, 12, 128, 131, 134,
 241, 249
 Defintion, **76**
 PMIX_JOB_TERM_STATUS
 Defintion, **79**
 PMIX_JOBID, 75, 127, 130, 134, 182, 187,
 241, 249
 Defintion, **73**
 pmix_key_t, 26, 58, 124, 126, 358
 Defintion, **26**
 PMIX_KVAL, 67
 PMIX_LAUNCH_COMPLETE, 23
 PMIX_LAUNCH_DIRECTIVE, 23
 PMIX_LAUNCHER
 Defintion, **70**
 PMIX_LAUNCHER_READY, 23
 PMIX_LOCAL, 33
 PMIX_LOCAL_CPUSSETS, 242, 252
 Defintion, **74**
 PMIX_LOCAL_PEERS, 242, 252
 Defintion, **74**
 PMIX_LOCAL_PROCS, 244
 Defintion, **74**

PMIX_LOCAL_RANK, 182, 183, 187, 188,
 242
 Defintion, **73**
 PMIX_LOCAL_SIZE, 242
 Defintion, **76**
 PMIX_LOCAL_TOPO
 Defintion, **77**
 PMIX_LOCALITY
 Defintion, **74**
 PMIX_LOCALITY_STRING
 Defintion, **77**
 PMIX_LOCALLDR, 244
 Defintion, **74**
 PMIx_Log, 10, 210, 385
 Defintion, **208**
 PMIX_LOG_EMAIL, 210, 213, 302
 Defintion, **86**
 PMIX_LOG_EMAIL_ADDR, 210, 213, 302
 Defintion, **86**
 PMIX_LOG_EMAIL_MSG, 210, 213, 302
 Defintion, **86**
 PMIX_LOG_EMAIL_SENDER_ADDR
 Defintion, **86**
 PMIX_LOG_EMAIL_SERVER
 Defintion, **86**
 PMIX_LOG_EMAIL_SRVR_PORT
 Defintion, **86**
 PMIX_LOG_EMAIL_SUBJECT, 210, 213,
 302
 Defintion, **86**
 PMIX_LOG_GENERATE_TIMESTAMP,
 210, 212
 Defintion, **86**
 PMIX_LOG_GLOBAL_DATASTORE, 210,
 213
 Defintion, **86**
 PMIX_LOG_GLOBAL_SYSLOG, 209, 212
 Defintion, **85**
 PMIX_LOG_JOB_RECORD, 210, 213
 Defintion, **86**
 PMIX_LOG_LOCAL_SYSLOG, 209, 212
 Defintion, **85**
 PMIX_LOG_MSG, 302
 Defintion, **86**
 PMIx_Log_nb, 9, 85, 213
 Defintion, **210**
 PMIX_LOG_ONCE, 209, 212
 Defintion, **86**
 PMIX_LOG_SOURCE, 209, 212
 Defintion, **85**
 PMIX_LOG_STDERR, 209, 212, 301
 Defintion, **85**
 PMIX_LOG_STDOUT, 209, 212, 301
 Defintion, **85**
 PMIX_LOG_SYSLOG, 209, 212, 301
 Defintion, **85**
 PMIX_LOG_SYSLOG_PRI, 209, 212
 Defintion, **85**
 PMIX_LOG_TAG_OUTPUT, 210, 212
 Defintion, **86**
 PMIX_LOG_TIMESTAMP, 209, 212
 Defintion, **85**
 PMIX_LOG_TIMESTAMP_OUTPUT, 210,
 212
 Defintion, **86**
 PMIX_LOG_XML_OUTPUT, 210, 212
 Defintion, **86**
 PMIx_Lookup, 8, 49, 142, 148, 149, 380,
 381
 Defintion, **145**
 pmix_lookup_cbfunc_t, 96, 280
 Defintion, **96**
 PMIx_Lookup_nb, 96, 97
 Defintion, **148**
 PMIX_MAP_BLOB
 Defintion, **80**
 PMIX_MAPBY, 156, 161, 243, 286
 Defintion, **82**
 PMIX_MAPPER, 82, 156, 161, 285
 Defintion, **82**
 PMIX_MAX_KEYLEN, 20
 PMIX_MAX_NSLEN, 20
 PMIX_MAX_PROCS, 12, 56, 76, 77, 136,
 241, 242
 Defintion, **77**
 PMIX_MAX_RESTARTS, 158, 163, 287

Definition, [84](#)
 PMIX_MAX_VALUE, [12](#), [56](#)
 Definition, [92](#)
 PMIX_MERGE_STDERR_STDOUT, [157](#),
 [162](#), [286](#)
 Definition, [83](#)
 PMIX_MIN_VALUE, [12](#), [56](#)
 Definition, [92](#)
 PMIX_MODEL_AFFINITY_POLICY
 Definition, [71](#)
 PMIX_MODEL_CPU_TYPE
 Definition, [71](#)
 PMIX_MODEL_DECLARED, [23](#)
 PMIX_MODEL_LIBRARY_NAME
 Definition, [71](#)
 PMIX_MODEL_LIBRARY_VERSION
 Definition, [71](#)
 PMIX_MODEL_NUM_CPUS
 Definition, [71](#)
 PMIX_MODEL_NUM_THREADS
 Definition, [71](#)
 PMIX_MODEL_PHASE_NAME
 Definition, [71](#)
 PMIX_MODEL_PHASE_TYPE
 Definition, [71](#)
 PMIX_MODEL_RESOURCES, [23](#)
 pmix_modex_cbfunc_t, [94](#), [273](#), [276](#)
 Definition, [94](#)
 PMIX_MONITOR_APP_CONTROL, [205](#),
 [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_CANCEL, [204](#), [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_FILE, [205](#), [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_FILE_ACCESS, [205](#),
 [207](#), [310](#)
 Definition, [91](#)
 PMIX_MONITOR_FILE_ALERT, [23](#)
 PMIX_MONITOR_FILE_CHECK_TIME,
 [205](#), [207](#), [310](#)
 Definition, [91](#)
 PMIX_MONITOR_FILE_DROPS, [205](#),
 [207](#), [310](#)
 Definition, [91](#)
 PMIX_MONITOR_FILE_MODIFY, [205](#),
 [207](#), [310](#)
 Definition, [91](#)
 PMIX_MONITOR_FILE_SIZE, [205](#), [207](#),
 [310](#)
 Definition, [91](#)
 PMIX_MONITOR_HEARTBEAT, [205](#),
 [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_HEARTBEAT_ALERT,
 [23](#)
 PMIX_MONITOR_HEARTBEAT_DROPS,
 [205](#), [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_HEARTBEAT_TIME,
 [205](#), [207](#), [310](#)
 Definition, [90](#)
 PMIX_MONITOR_ID, [204](#), [206](#), [310](#)
 Definition, [90](#)
 PMIX_NETWORK_COORDINATE
 Definition, [355](#)
 PMIX_NETWORK_COORDS_UPDATED,
 [355](#)
 PMIX_NETWORK_DIMS
 Definition, [355](#)
 PMIX_NETWORK_ENDPT
 Definition, [356](#)
 PMIX_NETWORK_NIC
 Definition, [355](#)
 PMIX_NETWORK_PLANE, [321](#), [353](#)
 Definition, [355](#)
 PMIX_NETWORK_SHAPE
 Definition, [356](#)
 PMIX_NETWORK_VIEW
 Definition, [355](#)
 PMIX_NO_OVERSUBSCRIBE, [158](#), [162](#),
 [287](#)
 Definition, [83](#)
 PMIX_NO_PROCS_ON_HEAD, [158](#), [162](#),
 [287](#)
 Definition, [83](#)

PMIX_NODE_INFO, 127, 130, 136, 182,
 187
 Defintion, **75**
 PMIX_NODE_INFO_ARRAY, 76, 249, 251
 Defintion, **76**
 PMIX_NODE_LIST
 Defintion, **74**
 PMIX_NODE_MAP, 12, 74, 242, 248–250,
 261
 Defintion, **80**
 PMIX_NODE_RANK, 242
 Defintion, **73**
 PMIX_NODE_SIZE, 136, 244
 Defintion, **77**
 PMIX_NODEID, 75, 76, 127, 130, 136, 182,
 187, 242
 Defintion, **74**
 PMIX_NON_PMI, 157, 161, 286
 Defintion, **83**
 pmix_notification_fn_t, 100, 217, 364
 Defintion, **100**
 PMIX_NOTIFY_ALLOC_COMPLETE, 22
 PMIX_NOTIFY_COMPLETION, 158
 Defintion, **79**
 PMIx_Notify_event, 9, 391
 Defintion, **220**
 PMIX_NPROC_OFFSET, 243
 Defintion, **73**
 PMIX_NSDIR, 73
 Defintion, **73**
 PMIX_NSPACE, 75, 127, 130, 134, 182,
 183, 187, 188, 249
 Defintion, **73**
 pmix_nspace_t, 27, 30, 95, 358, 359
 Defintion, **27**
 PMIX_NUM_NODES, 128, 131, 132, 134,
 248, 249
 Defintion, **77**
 PMIX_NUM_SLOTS
 Defintion, **77**
 pmix_op_cbfunc_t, 96, 99, 102, 144, 151,
 167, 171, 176, 211, 220, 221, 241,
 254–256, 263, 264, 266, 269, 270,
 272, 277, 282, 288, 290, 292, 294,
 295, 301, 306, 309, 315, 317, 337,
 350
 Defintion, **96**
 PMIX_OPENMP_PARALLEL_ENTERED,
 23
 PMIX_OPENMP_PARALLEL_EXITED,
 23
 PMIX_OPERATION_IN_PROGRESS, 23
 PMIX_OPERATION_SUCCEEDED, 23
 PMIX_OPTIONAL, 126, 129
 Defintion, **79**
 PMIX_OUTPUT_TO_FILE, 157, 162, 287
 Defintion, **83**
 PMIX_PARENT_ID, 155, 160, 285
 Defintion, **74**
 PMIX_PDATA, 67
 PMIX_PDATA_CONSTRUCT
 Defintion, **50**
 PMIX_PDATA_CREATE
 Defintion, **50**
 PMIX_PDATA_DESTRUCT
 Defintion, **50**
 PMIX_PDATA_FREE
 Defintion, **50**
 PMIX_PDATA_LOAD
 Defintion, **51**
 pmix_pdata_t, 49–52, 96, 97, 147, 360
 Defintion, **49**
 PMIX_PDATA_XFER
 Defintion, **52**
 PMIX_PERSIST, 68
 PMIX_PERSIST_APP, 34
 PMIX_PERSIST_FIRST_READ, 34
 PMIX_PERSIST_INDEF, 34
 PMIX_PERSIST_INVALID, 34
 PMIX_PERSIST_PROC, 34
 PMIX_PERSIST_SESSION, 34
 PMIX_PERSISTENCE, 143, 145, 278
 Defintion, **79**
 PMIx_Persistence_string, 9, 393
 Defintion, **109**
 pmix_persistence_t, 34, 68, 109, 359, 393

Defintion, [34](#)
 PMIX_PERSONALITY, [156](#), [161](#), [285](#)
 Defintion, [82](#)
 PMIX_PID, [67](#)
 PMIX_POINTER, [68](#)
 PMIX_PPR, [156](#), [161](#), [286](#)
 Defintion, [82](#)
 PMIX_PREFIX, [156](#), [160](#), [285](#)
 Defintion, [82](#)
 PMIX_PRELOAD_BIN, [156](#), [161](#), [285](#)
 Defintion, [83](#)
 PMIX_PRELOAD_FILES, [156](#), [161](#), [285](#)
 Defintion, [83](#)
 PMIX_PREPEND_ENVAR
 Defintion, [87](#)
 PMIX_PROC, [67](#)
 PMIX_PROC_BLOB
 Defintion, [80](#)
 PMIX_PROC_CONSTRUCT, [29](#)
 Defintion, [59](#)
 PMIX_PROC_CREATE
 Defintion, [29](#)
 PMIX_PROC_DATA, [250](#)
 Defintion, [80](#)
 PMIX_PROC_DESTRUCT
 Defintion, [29](#)
 PMIX_PROC_FREE, [180](#)
 Defintion, [29](#)
 PMIX_PROC_HAS_CONNECTED, [23](#)
 PMIX_PROC_INFO, [68](#)
 PMIX_PROC_INFO_CONSTRUCT
 Defintion, [32](#)
 PMIX_PROC_INFO_CREATE
 Defintion, [33](#)
 PMIX_PROC_INFO_DESTRUCT
 Defintion, [32](#)
 PMIX_PROC_INFO_FREE
 Defintion, [33](#)
 pmix_proc_info_t, [31–33](#), [68](#), [84](#), [184](#), [189](#),
 [298](#), [359](#)
 Defintion, [31](#)
 PMIX_PROC_LOAD
 Defintion, [30](#)
 PMIX_PROC_MAP, [12](#), [242](#), [248](#), [249](#), [261](#)
 Defintion, [80](#)
 PMIX_PROC_PID
 Defintion, [74](#)
 PMIX_PROC_RANK, [68](#)
 PMIX_PROC_STATE, [68](#)
 PMIX_PROC_STATE_ABORTED, [31](#)
 PMIX_PROC_STATE_ABORTED_BY_SIG,
 [31](#)
 PMIX_PROC_STATE_CALLED_ABORT,
 [31](#)
 PMIX_PROC_STATE_CANNOT_RESTART,
 [31](#)
 PMIX_PROC_STATE_COMM_FAILED,
 [31](#)
 PMIX_PROC_STATE_CONNECTED, [31](#)
 PMIX_PROC_STATE_ERROR, [31](#)
 PMIX_PROC_STATE_FAILED_TO_LAUNCH,
 [31](#)
 PMIX_PROC_STATE_FAILED_TO_START,
 [31](#)
 PMIX_PROC_STATE_KILLED_BY_CMD,
 [31](#)
 PMIX_PROC_STATE_LAUNCH_UNDERWAY,
 [31](#)
 PMIX_PROC_STATE_MIGRATING, [31](#)
 PMIX_PROC_STATE_PREPPED, [31](#)
 PMIX_PROC_STATE_RESTART, [31](#)
 PMIX_PROC_STATE_RUNNING, [31](#)
 PMIX_PROC_STATE_STATUS
 Defintion, [79](#)
 PMIx_Proc_state_string, [9](#), [392](#)
 Defintion, [108](#)
 pmix_proc_state_t, [30](#), [68](#), [108](#), [359](#), [392](#)
 Defintion, [30](#)
 PMIX_PROC_STATE_TERM_NON_ZERO,
 [31](#)
 PMIX_PROC_STATE_TERM_WO_SYNC,
 [31](#)
 PMIX_PROC_STATE_TERMINATE, [31](#)
 PMIX_PROC_STATE_TERMINATED, [31](#)
 PMIX_PROC_STATE_UNDEF, [31](#)
 PMIX_PROC_STATE_UNTERMINATED,

31
 pmix_proc_t, 28–30, 51, 58, 67, 74, 81, 100,
 104, 107, 114, 116, 118, 120, 128,
 138–140, 153, 218, 221, 222, 227,
 228, 244, 255–258, 264, 269, 270,
 272, 273, 276, 277, 280, 282, 284,
 288, 290, 295, 297, 301, 303, 306,
 308, 311, 313, 315, 317, 328, 332,
 339, 342, 346, 359
 Definition, **28**
 PMIX_PROC_TERM_STATUS
 Definition, **79**
 PMIX_PROC_TERMINATED, 23
 PMIX_PROC_URI, 185, 189
 Definition, **74**
 PMIX_PROCDIR
 Definition, **73**
 PMIx_Process_monitor, 11, 207
 Definition, **204**
 PMIx_Process_monitor_nb, 9, 90, 179, 208,
 386
 Definition, **206**
 PMIX_PROCID, 182, 183, 187, 188, 244
 Definition, **73**
 PMIX_PROGRAMMING_MODEL
 Definition, **71**
 PMIX_PSET_NAME, 324
 Definition, **70**
 PMIx_Publish, 8, 34, 79, 143–145, 278, 279,
 380, 381
 Definition, **142**
 PMIx_Publish_nb, 8, 145
 Definition, **144**
 PMIx_Put, 8, 33, 34, 36, 103, 124, 125, 128,
 131, 137, 139, 163, 190, 258, 277,
 331, 341, 378
 Definition, **124**
 PMIX_QUERY, 68
 PMIX_QUERY_ALLOC_STATUS, 184,
 189, 299
 Definition, **85**
 PMIX_QUERY_ATTRIBUTE_SUPPORT,
 182, 187, 190
 Definition, **85**
 PMIX_QUERY_AUTHORIZATIONS
 Definition, **84**
 PMIX_QUERY_CONSTRUCT
 Definition, **54**
 PMIX_QUERY_CREATE
 Definition, **55**
 PMIX_QUERY_DEBUG_SUPPORT, 184,
 189, 298
 Definition, **85**
 PMIX_QUERY_DESTRUCT
 Definition, **55**
 PMIX_QUERY_FREE
 Definition, **55**
 PMIx_Query_info, 185, 190
 Definition, **181**
 PMIx_Query_info_nb, 9, 10, 54, 76, 84, 136,
 163, 179, 185, 190, 191, 262, 324,
 325, 353, 384
 Definition, **185**
 PMIX_QUERY_JOB_STATUS, 184, 189,
 298
 Definition, **84**
 PMIX_QUERY_LOCAL_ONLY, 298
 Definition, **85**
 PMIX_QUERY_LOCAL_PROC_TABLE,
 184, 189, 298
 Definition, **84**
 PMIX_QUERY_MEMORY_USAGE, 184,
 189, 298
 Definition, **85**
 PMIX_QUERY_NAMESPACES, 184, 188,
 298
 Definition, **84**
 PMIX_QUERY_NUM_PSETS
 Definition, **85**
 PMIX_QUERY_PARTIAL_SUCCESS, 22
 PMIX_QUERY_PROC_TABLE, 184, 189,
 298
 Definition, **84**
 PMIX_QUERY_PSET_NAMES
 Definition, **85**
 PMIX_QUERY_QUALIFIERS_CREATE,

10, 11
 Definition, **55**
 PMIX_QUERY_QUEUE_LIST, 184, 189,
 298
 Definition, **84**
 PMIX_QUERY_QUEUE_STATUS, 184,
 189, 298
 Definition, **84**
 PMIX_QUERY_REFRESH_CACHE, 181,
 185, 186, 190
 Definition, **84**
 PMIX_QUERY_REPORT_AVG, 184, 189,
 298
 Definition, **85**
 PMIX_QUERY_REPORT_MINMAX, 184,
 189, 299
 Definition, **85**
 PMIX_QUERY_SPAWN_SUPPORT, 184,
 189, 298
 Definition, **84**
 pmix_query_t, 10, 11, 54, 55, 68, 183, 188,
 190, 297, 299, 361
 Definition, **54**
 PMIX_RANGE, 143, 145, 146, 149, 150,
 152, 205, 218, 278, 280, 283, 296
 Definition, **79**
 PMIX_RANGE_CUSTOM, 34
 PMIX_RANGE_GLOBAL, 34
 PMIX_RANGE_INVALID, 34
 PMIX_RANGE_LOCAL, 34
 PMIX_RANGE_NAMESPACE, 34
 PMIX_RANGE_PROC_LOCAL, 34
 PMIX_RANGE_RM, 34
 PMIX_RANGE_SESSION, 34
 PMIX_RANGE_UNDEF, 34
 PMIX_RANK, 182, 183, 187, 188, 242
 Definition, **73**
 PMIX_RANK_INVALID, 28
 PMIX_RANK_LOCAL_NODE, 28
 PMIX_RANK_LOCAL_PEERS, 28
 pmix_rank_t, 28, 30, 68, 359
 Definition, **28**
 PMIX_RANK_UNDEF, 28
 PMIX_RANK_VALID, 28
 PMIX_RANK_WILDCARD, 28
 PMIX_RANKBY, 157, 161, 243, 286
 Definition, **82**
 PMIX_RECONNECT_SERVER
 Definition, **70**
 PMIX_REGATTR, 68
 PMIX_REGATTR_CONSTRUCT
 Definition, **57**
 PMIX_REGATTR_CREATE
 Definition, **57**
 PMIX_REGATTR_DESTRUCT
 Definition, **57**
 PMIX_REGATTR_FREE
 Definition, **58**
 PMIX_REGATTR_LOAD
 Definition, **58**
 pmix_regattr_t, 12, 56–58, 68, 92, 191, 261,
 262, 361
 Definition, **56**
 PMIX_REGATTR_XFER
 Definition, **58**
 PMIX_REGEX, 68
 PMIx_Register_attributes, 12, 401
 Definition, **261**
 PMIX_REGISTER_CLEANUP, 199, 202
 Definition, **90**
 PMIX_REGISTER_CLEANUP_DIR, 199,
 202
 Definition, **90**
 PMIx_Register_event_handler, 9, 99, 179,
 390
 Definition, **216**
 PMIX_REGISTER_NODATA, 241
 Definition, **79**
 pmix_release_cbfunc_t, 94
 Definition, **94**
 PMIX_REMOTE, 33
 PMIX_REPORT_BINDINGS, 158, 162, 287
 Definition, **84**
 PMIX_REQUESTOR_IS_CLIENT, 155,
 160
 Definition, **70**

PMIX_REQUESTOR_IS_TOOL, 155, 160
 Defintion, **70**
 PMIx_Resolve_nodes, 8, 384
 Defintion, **180**
 PMIx_Resolve_peers, 8, 383
 Defintion, **179**
 PMIX_RM_NAME
 Defintion, **87**
 PMIX_RM_VERSION
 Defintion, **87**
 PMIX_SCOPE, 68
 PMIx_Scope_string, 9, 393
 Defintion, **109**
 pmix_scope_t, 33, 68, 109, 125, 359, 392
 Defintion, **33**
 PMIX_SCOPE_UNDEF, 33
 PMIX_SEND_HEARTBEAT
 Defintion, **90**
 pmix_server_abort_fn_t, 365
 Defintion, **271**
 pmix_server_alloc_fn_t, 372
 Defintion, **302**
 PMIX_SERVER_ATTRIBUTES, 12, 182, 187
 Defintion, **92**
 pmix_server_client_connected_fn_t, 96, 233, 256, 269, 364
 Defintion, **268**
 pmix_server_client_finalized_fn_t, 271, 364
 Defintion, **270**
 PMIx_server_collect_inventory, 11, 403
 Defintion, **265**
 pmix_server_connect_fn_t, 163, 289, 291, 368
 Defintion, **288**
 PMIx_server_deliver_inventory, 11, 403
 Defintion, **266**
 PMIx_server_deregister_client, 8, 399
 Defintion, **256**
 pmix_server_deregister_events_fn_t, 370
 Defintion, **293**
 PMIx_server_deregister_fabric
 Defintion, **321**
 PMIx_server_deregister_nspace, 8, 256, 398
 Defintion, **253**
 pmix_server_disconnect_fn_t, 291, 369
 Defintion, **290**
 pmix_server_dmodex_req_fn_t, 10, 11, 94, 366
 Defintion, **276**
 PMIx_server_dmodex_request, 8, 102, 103, 258, 400
 Defintion, **257**
 PMIX_SERVER_ENABLE_MONITORING
 Defintion, **69**
 pmix_server_fencebn_fn_t, 12, 94, 275, 365
 Defintion, **272**
 PMIx_server_finalize, 8, 396
 Defintion, **123**
 PMIX_SERVER_FUNCTIONS
 Defintion, **92**
 PMIX_SERVER_GATEWAY
 Defintion, **69**
 pmix_server_get_cred_fn_t, 314, 374
 Defintion, **311**
 PMIx_server_get_index
 Defintion, **322**
 PMIx_server_get_vertex_info
 Defintion, **322**
 PMIX_SERVER_HOSTNAME
 Defintion, **70**
 PMIx_server_init, 8, 112, 262, 267, 396
 Defintion, **121**
 PMIx_server_IOF_deliver, 11, 172, 402
 Defintion, **264**
 pmix_server_iof_fn_t, 375
 Defintion, **314**
 pmix_server_job_control_fn_t, 373
 Defintion, **305**
 pmix_server_listener_fn_t
 Defintion, **296**
 pmix_server_log_fn_t, 372
 Defintion, **300**
 pmix_server_lookup_fn_t, 367
 Defintion, **279**
 pmix_server_module_t, 121, 123, 262, 267

Defintion, [267](#)
 pmix_server_monitor_fn_t, [373](#)
 Defintion, [308](#)
 pmix_server_notify_event_fn_t, [101](#), [370](#)
 Defintion, [295](#)
 PMIX_SERVER_NAMESPACE, [120](#), [121](#), [243](#)
 Defintion, [69](#)
 PMIX_SERVER_PIDINFO, [116](#), [118](#), [120](#)
 Defintion, [69](#)
 pmix_server_publish_fn_t, [366](#)
 Defintion, [277](#)
 pmix_server_query_fn_t, [371](#)
 Defintion, [297](#)
 PMIX_SERVER_RANK, [121](#), [243](#)
 Defintion, [69](#)
 PMix_server_register_client, [8](#), [233](#), [256](#),
 [269](#), [271](#), [399](#)
 Defintion, [255](#)
 pmix_server_register_events_fn_t, [369](#)
 Defintion, [291](#)
 PMix_server_register_fabric
 Defintion, [320](#)
 PMix_server_register_namespace, [8](#), [10](#), [17](#), [75](#),
 [96](#), [239](#), [246](#), [249](#), [398](#)
 Defintion, [240](#)
 PMIX_SERVER_REMOTE_CONNECTIONS,
 [122](#)
 Defintion, [69](#)
 PMIX_SERVER_SCHEDULER, [319](#)
 Defintion, [320](#)
 PMix_server_setup_application, [9](#), [12](#), [101](#),
 [102](#), [263](#), [267](#), [401](#)
 Defintion, [259](#)
 PMix_server_setup_fork, [8](#), [400](#)
 Defintion, [257](#)
 PMix_server_setup_local_support, [9](#), [402](#)
 Defintion, [262](#)
 pmix_server_spawn_fn_t, [95](#), [368](#)
 Defintion, [283](#)
 pmix_server_stdin_fn_t, [375](#)
 Defintion, [317](#)
 PMIX_SERVER_SYSTEM_SUPPORT, [121](#)
 Defintion, [69](#)
 PMIX_SERVER_TMPDIR, [121](#)
 Defintion, [69](#)
 pmix_server_tool_connection_fn_t, [233](#), [371](#)
 Defintion, [299](#)
 PMIX_SERVER_TOOL_SUPPORT, [121](#),
 [123](#), [233](#)
 Defintion, [69](#)
 pmix_server_unpublish_fn_t, [367](#)
 Defintion, [281](#)
 PMIX_SERVER_URI, [116](#), [118](#), [120](#), [184](#),
 [189](#)
 Defintion, [70](#)
 pmix_server_validate_cred_fn_t, [374](#)
 Defintion, [312](#)
 PMIX_SESSION_ID, [75](#), [126](#), [130](#), [133](#),
 [182](#), [186](#), [243](#), [248](#), [249](#)
 Defintion, [74](#)
 PMIX_SESSION_INFO, [126](#), [130](#), [132](#), [181](#),
 [186](#)
 Defintion, [74](#)
 PMIX_SESSION_INFO_ARRAY, [10](#), [76](#),
 [241](#), [248](#)
 Defintion, [75](#)
 PMIX_SET_ENVAR
 Defintion, [87](#)
 PMIX_SET_SESSION_CWD, [156](#), [160](#), [285](#)
 Defintion, [83](#)
 PMIX_SETENV
 Defintion, [66](#)
 PMIX_SETUP_APP_ALL, [259](#)
 Defintion, [92](#)
 PMIX_SETUP_APP_ENVARS, [259](#)
 Defintion, [92](#)
 PMIX_SETUP_APP_NONENVARS, [259](#)
 Defintion, [92](#)
 pmix_setup_application_cbfunc_t, [259](#)
 Defintion, [101](#)
 PMIX_SINGLE_LISTENER, [113](#)
 Defintion, [71](#)
 PMIX_SIZE, [67](#)
 PMIX_SOCKET_MODE, [113](#), [117](#), [122](#)
 Defintion, [71](#)
 PMix_Spawn, [8](#), [52](#), [73](#), [82](#), [86](#), [87](#), [154](#),

155, 159, 160, 163, 245, 257, 283,
 288, 367, 382
 Definition, **154**
 pmix_spawn_cbfunc_t, 95, 159, 284
 Definition, **95**
 PMIx_Spawn_nb, 8, 52, 95
 Definition, **159**
 PMIX_SPAWN_TOOL
 Definition, **84**
 PMIX_SPAWNED, 155, 160, 285
 Definition, **72**
 PMIX_STATUS, 67
 pmix_status_t, 21, 25, 39–41, 62, 63, 66, 67,
 98–100, 102–106, 108, 216, 221,
 292, 294, 295, 358, 370, 392
 Definition, **21**
 PMIX_STDIN_TGT, 157, 161, 286
 Definition, **83**
 PMIx_Store_internal, 8, 378
 Definition, **131**
 PMIX_STRING, 67
 PMIX_SUCCESS, 21
 pmix_system_event
 Definition, **25**
 PMIX_SYSTEM_TMPDIR, 121
 Definition, **69**
 PMIX_TAG_OUTPUT, 157, 162, 286
 Definition, **83**
 PMIX_TCP_DISABLE_IPV4, 114, 117, 122
 Definition, **72**
 PMIX_TCP_DISABLE_IPV6, 114, 117, 122
 Definition, **72**
 PMIX_TCP_IF_EXCLUDE, 113, 117, 122
 Definition, **72**
 PMIX_TCP_IF_INCLUDE, 113, 117, 122
 Definition, **72**
 PMIX_TCP_IPV4_PORT, 114, 117, 122
 Definition, **72**
 PMIX_TCP_IPV6_PORT, 114, 117, 122
 Definition, **72**
 PMIX_TCP_REPORT_URI, 113, 117, 122
 Definition, **72**
 PMIX_TCP_URI, 116, 118
 Definition, **72**
 PMIX_TDIR_RMCLEAN
 Definition, **73**
 PMIX_THREADING_MODEL
 Definition, **71**
 PMIX_TIME, 67
 PMIX_TIME_REMAINING, 179, 184, 189,
 299
 Definition, **85**
 PMIX_TIMEOUT, 3, 13, 127, 128, 130, 131,
 138, 139, 141, 143, 145–147, 149,
 150, 152, 165, 168, 169, 171, 172,
 235, 237, 274, 277, 278, 281, 283,
 287, 289, 291, 312, 314, 327, 329,
 331, 334–336, 338, 340, 344, 346,
 348, 349
 Definition, **78**
 PMIX_TIMESTAMP_OUTPUT, 157, 162,
 286
 Definition, **83**
 PMIX_TIMEVAL, 67
 PMIX_TMPDIR, 73
 Definition, **73**
 PMIX_TOOL_ATTRIBUTES, 12, 183, 188
 Definition, **92**
 PMIx_tool_connect_to_server, 11, 405
 Definition, **119**
 pmix_tool_connection_cbfunc_t, 299
 Definition, **103**
 PMIX_TOOL_DO_NOT_CONNECT, 116,
 118
 Definition, **70**
 PMIx_tool_finalize, 9, 404
 Definition, **119**
 PMIX_TOOL_FUNCTIONS
 Definition, **92**
 PMIx_tool_init, 9, 69, 112, 119, 404
 Definition, **115**
 PMIX_TOOL_NAMESPACE, 116
 Definition, **69**
 PMIX_TOOL_RANK, 116
 Definition, **69**
 PMIX_TOPOLOGY

Defintion, [77](#)
 PMIX_TOPOLOGY_FILE
 Defintion, [77](#)
 PMIX_TOPOLOGY_SIGNATURE
 Defintion, [77](#)
 PMIX_TOPOLOGY_XML
 Defintion, [77](#)
 PMIX_UINT, [67](#)
 PMIX_UINT16, [67](#)
 PMIX_UINT32, [67](#)
 PMIX_UINT64, [67](#)
 PMIX_UINT8, [67](#)
 PMIX_UNDEF, [67](#)
 PMIX_UNIV_SIZE, [10](#), [12](#), [128](#), [131](#), [132](#),
 [241](#), [248](#)
 Defintion, [76](#)
 PMIx_Unpublish, [8](#), [151](#), [152](#), [381](#)
 Defintion, [149](#)
 PMIx_Unpublish_nb, [8](#)
 Defintion, [151](#)
 PMIX_UNSET_ENVAR
 Defintion, [87](#)
 PMIX_USERID, [106](#), [142](#), [144](#), [146](#), [148](#),
 [150](#), [152](#), [183](#), [188](#), [192](#), [195](#), [198](#),
 [201](#), [204](#), [206](#), [209](#), [211](#), [235](#), [237](#),
 [278–284](#), [292](#), [298](#), [300](#), [301](#), [303](#),
 [306](#), [309](#), [311](#), [313](#), [315](#), [318](#)
 Defintion, [70](#)
 PMIX_USOCK_DISABLE, [113](#), [122](#)
 Defintion, [71](#)
 PMIx_Validate_credential, [11](#), [387](#)
 Defintion, [236](#)
 pmix_validation_cbfunc_t, [236](#), [313](#)
 Defintion, [105](#)
 PMIX_VALUE, [67](#)
 pmix_value_cbfunc_t, [17](#), [97](#)
 Defintion, [97](#)
 PMIX_VALUE_CONSTRUCT
 Defintion, [38](#)
 PMIX_VALUE_CREATE
 Defintion, [38](#)
 PMIX_VALUE_DESTRUCT
 Defintion, [38](#)
 PMIX_VALUE_FREE
 Defintion, [38](#)
 PMIX_VALUE_GET_NUMBER
 Defintion, [41](#)
 PMIX_VALUE_LOAD
 Defintion, [39](#)
 pmix_value_t, [36–41](#), [67](#), [97](#), [124](#), [125](#), [322](#),
 [323](#), [360](#)
 Defintion, [36](#)
 PMIX_VALUE_UNLOAD
 Defintion, [39](#)
 PMIX_VALUE_XFER
 Defintion, [40](#)
 PMIX_VERSION_INFO
 Defintion, [70](#)
 PMIX_WAIT, [146](#), [147](#), [149](#), [280](#)
 Defintion, [78](#)
 PMIX_WDIR, [155](#), [160](#), [285](#)
 Defintion, [82](#)
 PMIxClient.abort
 Defintion, [377](#)
 PMIxClient.alloc_directive_string
 Defintion, [394](#)
 PMIxClient.allocate
 Defintion, [385](#)
 PMIxClient.commit
 Defintion, [378](#)
 PMIxClient.connect
 Defintion, [382](#)
 PMIxClient.data_range_string
 Defintion, [393](#)
 PMIxClient.data_type_string
 Defintion, [394](#)
 PMIxClient.deregister_event_handler
 Defintion, [391](#)
 PMIxClient.disconnect
 Defintion, [382](#)
 PMIxClient.error_string
 Defintion, [392](#)
 PMIxClient.fence
 Defintion, [379](#)
 PMIxClient.finalize
 Defintion, [377](#)

PMIxClient.get
 Defintion, [379](#)
 PMIxClient.get_credential
 Defintion, [386](#)
 PMIxClient.get_version
 Defintion, [376](#)
 PMIxClient.group_construct
 Defintion, [387](#)
 PMIxClient.group_destruct
 Defintion, [390](#)
 PMIxClient.group_invite
 Defintion, [388](#)
 PMIxClient.group_join
 Defintion, [389](#)
 PMIxClient.group_leave
 Defintion, [389](#)
 PMIxClient.info_directives_string
 Defintion, [394](#)
 PMIxClient.init
 Defintion, [376](#)
 PMIxClient.initialized
 Defintion, [376](#)
 PMIxClient.iof_channel_string
 Defintion, [395](#)
 PMIxClient.job_ctrl
 Defintion, [385](#)
 PMIxClient.log
 Defintion, [384](#)
 PMIxClient.lookup
 Defintion, [380](#)
 PMIxClient.monitor
 Defintion, [386](#)
 PMIxClient.notify_event
 Defintion, [391](#)
 PMIxClient.persistance_string
 Defintion, [393](#)
 PMIxClient.proc_state_string
 Defintion, [392](#)
 PMIxClient.publish
 Defintion, [380](#)
 PMIxClient.put, [378](#)
 Defintion, [378](#)
 PMIxClient.query
 Defintion, [384](#)
 PMIxClient.register_event_handler
 Defintion, [390](#)
 PMIxClient.resolve_nodes
 Defintion, [383](#)
 PMIxClient.resolve_peers
 Defintion, [383](#)
 PMIxClient.scope_string
 Defintion, [392](#)
 PMIxClient.spawn
 Defintion, [381](#)
 PMIxClient.store_internal
 Defintion, [377](#)
 PMIxClient.unpublish
 Defintion, [381](#)
 PMIxClient.validate_credential
 Defintion, [387](#)
 PMIxServer.collect_inventory
 Defintion, [402](#)
 PMIxServer.deliver_inventory
 Defintion, [403](#)
 PMIxServer.deregister_client
 Defintion, [399](#)
 PMIxServer.deregister_namespace
 Defintion, [398](#)
 PMIxServer.dmodex_request
 Defintion, [400](#)
 PMIxServer.finalize
 Defintion, [396](#)
 PMIxServer.generate_ppn
 Defintion, [397](#)
 PMIxServer.generate_regex
 Defintion, [396](#)
 PMIxServer.init
 Defintion, [395](#)
 PMIxServer.iof_deliver
 Defintion, [402](#)
 PMIxServer.register_attributes
 Defintion, [401](#)
 PMIxServer.register_client
 Defintion, [398](#)
 PMIxServer.register_namespace
 Defintion, [397](#)

PMIxServer.setup_application
 Defintion, [400](#)

PMIxServer.setup_fork
 Defintion, [399](#)

PMIxServer.setup_local_support
 Defintion, [401](#)

PMIxTool.connect_to_server
 Defintion, [404](#)

PMIxTool.finalize
 Defintion, [404](#)

PMIxTool.init
 Defintion, [403](#)

PMIxTool.iof_deregister
 Defintion, [405](#)

PMIxTool.iof_pull
 Defintion, [405](#)

PMIxTool.iof_push
 Defintion, [406](#)

rank, [134](#), [250](#)
 Defintion, [14](#)

resource manager
 Defintion, [15](#)

scheduler, [319](#)
 Defintion, [15](#)

session, [9](#), [10](#), [12](#), [74](#), [75](#), [132](#), [190](#), [246](#)
 Defintion, [14](#)

slot
 Defintion, [14](#)

slots
 Defintion, [14](#)

workflow
 Defintion, [14](#)

workflows, [105](#)
 Defintion, [14](#)