



Process Management Interface for Exascale (PMIx) Standard

Version 2.1

December 2018

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 2.1.

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 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?	3
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	5
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	9
1.7. Version 2.1: Dec. 2018	10
2. PMIx Terms and Conventions	11
2.1. Notational Conventions	13
2.2. Semantics	14
2.3. Naming Conventions	14
2.4. Procedure Conventions	15
2.5. Standard vs Reference Implementation	15
3. Data Structures and Types	16
3.1. Constants	16
3.1.1. Error Constants	17
3.2. Data Types	20
3.2.1. Key Structure	20
3.2.2. Namespace Structure	21
3.2.3. Rank Structure	22

3.2.4.	Process Structure	22
3.2.5.	Process structure support macros	22
3.2.6.	Process State Structure	24
3.2.7.	Process Information Structure	25
3.2.8.	Process Information Structure support macros	26
3.2.9.	Scope of Put Data	27
3.2.10.	Range of Published Data	28
3.2.11.	Data Persistence Structure	28
3.2.12.	Value Structure	29
3.2.13.	Value structure support macros	30
3.2.14.	Load a <code>pmix_value_t</code> structure	31
3.2.15.	Info and Info Array Structures	32
3.2.16.	Info structure support macros	33
3.2.17.	Info Type Directives	36
3.2.18.	Info Directive support macros	36
3.2.19.	Job Allocation Directives	37
3.2.20.	Lookup Returned Data Structure	38
3.2.21.	Lookup data structure support macros	38
3.2.22.	Application Structure	41
3.2.23.	App structure support macros	41
3.2.24.	Query Structure	43
3.2.25.	Query structure support macros	43
3.2.26.	Modex Structure	44
3.2.27.	Modex data structure support macros	45
3.3.	Data Packing/Unpacking Types and Structures	46
3.3.1.	Byte Object Type	46
3.3.2.	Byte object support macros	46
3.3.3.	Data Buffer Type	48
3.3.4.	Data buffer support macros	48
3.3.5.	Data Array Structure	50
3.3.6.	Generalized Data Types Used for Packing/Unpacking	50
3.4.	Reserved attributes	52
3.4.1.	Initialization attributes	52

3.4.2.	Tool-related attributes	53
3.4.3.	Identification attributes	53
3.4.4.	UNIX socket rendezvous socket attributes	54
3.4.5.	TCP connection attributes	54
3.4.6.	Global Data Storage (GDS) attributes	55
3.4.7.	General process-level attributes	55
3.4.8.	Scratch directory attributes	56
3.4.9.	Relative Rank Descriptive Attributes	56
3.4.10.	Information retrieval attributes	57
3.4.11.	Information storage attributes	58
3.4.12.	Size information attributes	58
3.4.13.	Memory information attributes	59
3.4.14.	Topology information attributes	59
3.4.15.	Request-related attributes	60
3.4.16.	Server-to-PMIX library attributes	61
3.4.17.	Server-to-Client attributes	62
3.4.18.	Event handler registration and notification attributes	62
3.4.19.	Fault tolerance attributes	64
3.4.20.	Spawn attributes	64
3.4.21.	Query attributes	66
3.4.22.	Log attributes	67
3.4.23.	Debugger attributes	68
3.4.24.	Resource manager attributes	68
3.4.25.	Environment variable attributes	69
3.4.26.	Job Allocation attributes	69
3.4.27.	Job control attributes	70
3.4.28.	Monitoring attributes	71
3.5.	Callback Functions	71
3.5.1.	Release Callback Function	72
3.5.2.	Modex Callback Function	72
3.5.3.	Spawn Callback Function	73
3.5.4.	Op Callback Function	74
3.5.5.	Lookup Callback Function	74

3.5.6.	Value Callback Function	75
3.5.7.	Info Callback Function	75
3.5.8.	Event Handler Registration Callback Function	76
3.5.9.	Notification Handler Completion Callback Function	77
3.5.10.	Notification Function	78
3.5.11.	Server Setup Application Callback Function	80
3.5.12.	Server Direct Modex Response Callback Function	81
3.5.13.	<code>pmix_connection_cbfunc_t</code>	82
3.5.14.	<code>pmix_tool_connection_cbfunc_t</code>	83
3.5.15.	Constant String Functions	83
4.	Initialization and Finalization	86
4.1.	Query	86
4.1.1.	<code>PMIx_Initialized</code>	86
4.1.2.	<code>PMIx_Get_version</code>	87
4.2.	Client Initialization and Finalization	87
4.2.1.	<code>PMIx_Init</code>	87
4.2.2.	<code>PMIx_Finalize</code>	90
4.3.	Tool Initialization and Finalization	91
4.3.1.	<code>PMIx_tool_init</code>	91
4.3.2.	<code>PMIx_tool_finalize</code>	94
4.4.	Server Initialization and Finalization	95
4.4.1.	<code>PMIx_server_init</code>	95
4.4.2.	<code>PMIx_server_finalize</code>	97
5.	Key/Value Management	99
5.1.	Setting and Accessing Key/Value Pairs	99
5.1.1.	<code>PMIx_Put</code>	99
5.1.2.	<code>PMIx_Get</code>	100
5.1.3.	<code>PMIx_Get_nb</code>	103
5.1.4.	<code>PMIx_Store_internal</code>	106
5.1.5.	Accessing information: examples	107
5.2.	Exchanging Key/Value Pairs	112
5.2.1.	<code>PMIx_Commit</code>	112

5.2.2.	PMIx_Fence	112
5.2.3.	PMIx_Fence_nb	114
5.3.	Publish and Lookup Data	117
5.3.1.	PMIx_Publish	117
5.3.2.	PMIx_Publish_nb	119
5.3.3.	PMIx_Lookup	121
5.3.4.	PMIx_Lookup_nb	123
5.3.5.	PMIx_Unpublish	125
5.3.6.	PMIx_Unpublish_nb	127
6.	Process Management	129
6.1.	Abort	129
6.1.1.	PMIx_Abort	129
6.2.	Process Creation	130
6.2.1.	PMIx_Spawn	130
6.2.2.	PMIx_Spawn_nb	135
6.3.	Connecting and Disconnecting Processes	139
6.3.1.	PMIx_Connect	139
6.3.2.	PMIx_Connect_nb	142
6.3.3.	PMIx_Disconnect	144
6.3.4.	PMIx_Disconnect_nb	146
7.	Job Allocation Management and Reporting	148
7.1.	Query	148
7.1.1.	PMIx_Resolve_peers	149
7.1.2.	PMIx_Resolve_nodes	149
7.1.3.	PMIx_Query_info_nb	150
7.2.	Allocation Requests	155
7.2.1.	PMIx_Allocation_request_nb	156
7.2.2.	PMIx_Job_control_nb	158
7.3.	Process and Job Monitoring	161
7.3.1.	PMIx_Process_monitor_nb	161
7.3.2.	PMIx_Heartbeat	164

7.4.	Logging	164
7.4.1.	PMIx_Log_nb	164
8.	Event Notification	167
8.1.	Notification and Management	167
8.1.1.	PMIx_Register_event_handler	169
8.1.2.	PMIx_Deregister_event_handler	172
8.1.3.	PMIx_Notify_event	173
9.	Data Packing and Unpacking	176
9.1.	Support Macros	176
9.1.1.	PMIX_DATA_BUFFER_CREATE	176
9.1.2.	PMIX_DATA_BUFFER_RELEASE	177
9.1.3.	PMIX_DATA_BUFFER_CONSTRUCT	177
9.1.4.	PMIX_DATA_BUFFER_DESTRUCT	177
9.1.5.	PMIX_DATA_BUFFER_LOAD	178
9.1.6.	PMIX_DATA_BUFFER_UNLOAD	178
9.2.	General Routines	179
9.2.1.	PMIx_Data_pack	179
9.2.2.	PMIx_Data_unpack	181
9.2.3.	PMIx_Data_copy	183
9.2.4.	PMIx_Data_print	183
9.2.5.	PMIx_Data_copy_payload	184
10.	Server-Specific Interfaces	186
10.1.	Server Support Functions	186
10.1.1.	PMIx_generate_regex	186
10.1.2.	PMIx_generate_ppn	187
10.1.3.	PMIx_server_register_nspace	188
10.1.4.	PMIx_server_deregister_nspace	200
10.1.5.	PMIx_server_register_client	201
10.1.6.	PMIx_server_deregister_client	203
10.1.7.	PMIx_server_setup_fork	204
10.1.8.	PMIx_server_dmodex_request	204
10.1.9.	PMIx_server_setup_application	206

10.1.10. <code>PMIx_server_setup_local_support</code>	207
10.2. Server Function Pointers	209
10.2.1. <code>pmix_server_module_t</code> Module	209
10.2.2. <code>pmix_server_client_connected_fn_t</code>	210
10.2.3. <code>pmix_server_client_finalized_fn_t</code>	212
10.2.4. <code>pmix_server_abort_fn_t</code>	213
10.2.5. <code>pmix_server_fence_nb_fn_t</code>	214
10.2.6. <code>pmix_server_dmodex_req_fn_t</code>	217
10.2.7. <code>pmix_server_publish_fn_t</code>	218
10.2.8. <code>pmix_server_lookup_fn_t</code>	220
10.2.9. <code>pmix_server_unpublish_fn_t</code>	222
10.2.10. <code>pmix_server_spawn_fn_t</code>	224
10.2.11. <code>pmix_server_connect_fn_t</code>	229
10.2.12. <code>pmix_server_disconnect_fn_t</code>	231
10.2.13. <code>pmix_server_register_events_fn_t</code>	233
10.2.14. <code>pmix_server_deregister_events_fn_t</code>	235
10.2.15. <code>pmix_server_notify_event_fn_t</code>	237
10.2.16. <code>pmix_server_listener_fn_t</code>	238
10.2.17. <code>pmix_server_query_fn_t</code>	239
10.2.18. <code>pmix_server_tool_connection_fn_t</code>	242
10.2.19. <code>pmix_server_log_fn_t</code>	243
10.2.20. <code>pmix_server_alloc_fn_t</code>	245
10.2.21. <code>pmix_server_job_control_fn_t</code>	247
10.2.22. <code>pmix_server_monitor_fn_t</code>	250
A. Acknowledgements	253
A.1. Version 2.0	253
A.2. Version 1.0	254
Bibliography	256
Index	257

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 1.2.1 Who should use the standard?

2 The PMIx Standard informs PMIx clients and RMs of the syntax and semantics of the PMIx APIs.
3 PMIx clients (e.g., tools, Message Passing Environment (MPE) libraries) can use this standard to
4 understand the set of attributes provided by various APIs of the PRI and their intended behavior.
5 Additional information about the rationale for the selection of specific interfaces and attributes is
6 also provided.
7 PMIx-enabled RMs can use this standard to understand the expected behavior required of them
8 when they support various interfaces/attributes. In addition, optional features and suggestions on
9 behavior are also included in the discussion to help guide RM design and implementation.

10 1.2.2 What is defined in the standard?

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

15 1.2.3 What is *not* defined in the standard?

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

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

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

31 If an application (for example) truly relies on the [PMIX_TIMEOUT](#) attribute in a call to
32 [PMIx_Fence](#) , it should set the required flag in the [pmix_info_t](#) for that attribute. This
33 informs the library and its SMS host that it must return an immediate error if this attribute is not

1 supported. By not setting the flag, the library and SMS host are allowed to treat the attribute as
2 optional, ignoring it if support is not available.

3 It is therefore critical that users and application implementers:

- 4 a) consider whether or not a given attribute is required, marking it accordingly; and
- 5 b) check the return status on all PMIx function calls to ensure support was present and that the
6 request was accepted. Note that for non-blocking APIs, a return of **PMIX_SUCCESS** only
7 indicates that the request had no obvious errors and is being processed – the eventual callback
8 will return the status of the requested operation itself.

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

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

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

27 **1.2.4 General Guidance for PMIx Users and Implementors**

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

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

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

6 1.3 PMIx Architecture Overview

7 This section presents a brief overview of the PMIx Architecture [1]. Note that this is a conceptual
8 model solely used to help guide the standards process — it does not represent a design requirement
9 on any PMIx implementation. Instead, the model is used by the PMIx community as a sounding
10 board for evaluating proposed interfaces and avoid unintentionally imposing constraints on
11 implementers. Built into the model are two guiding principles also reflected in the standard. First,
12 PMIx operates in the mode of a *messenger*, and not a *doer* — i.e., the role of PMIx is to provide
13 communication between the various participants, relaying requests and returning responses. The
14 intent of the standard is not to suggest that PMIx itself actually perform any of the defined
15 operations — this is left to the various SMS elements and/or the application. Any exceptions to that
16 intent are left to the discretion of the particular implementation.

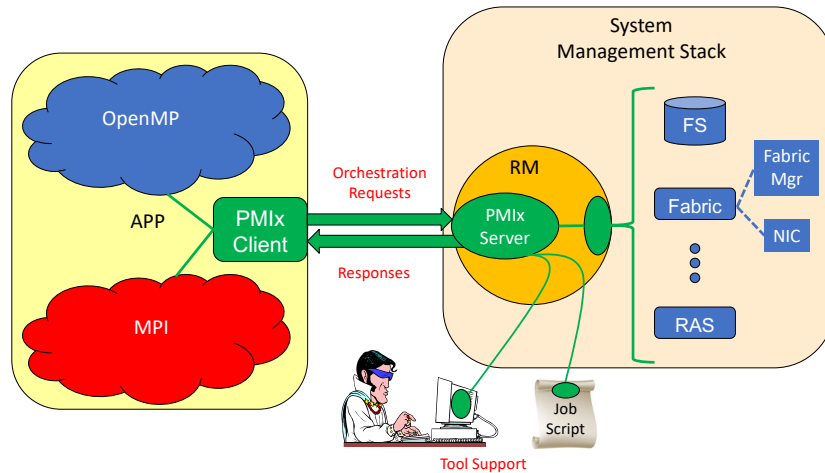


Figure 1.1.: PMIx-SMS Interactions

17 Thus, as the diagram in Fig. 1.1 shows, the application is built against a PMIx client library that
18 contains the client-side APIs, attribute definitions, and communication support for interacting with
19 the local PMIx server. Intra-process cross-library interactions are supported at the client level to
20 avoid unnecessary burdens on the server. Orchestration requests are sent to the local PMIx server,
21 which subsequently passes them to the host SMS (here represented by an RM daemon) using the

1 PMIx server callback functions the host SMS registered during `PMIx_server_init`. The host SMS
2 can indicate its lack of support for any operation by simply providing a `NULL` for the associated
3 callback function, or can create a function entry that returns *not supported* when called.

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

11 Recognizing the burden this places on SMS vendors, the PMIx community has included interfaces
12 by which the host can request support from local SMS elements. Once the SMS has transferred the
13 request to an appropriate location, a PMIx server interface can be used to pass the request between
14 SMS subsystems. For example, a request for network traffic statistics can utilize the PMIx
15 networking abstractions to retrieve the information from the FM. This reduces the portability and
16 interoperability issues between the individual subsystems by transferring the burden of defining the
17 interoperable interfaces from the SMS subsystems to the PMIx community, which continues to
18 work with those providers to develop the necessary support.

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

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

29 **1.3.1 The PMIx Reference Implementation (PRI)**

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

- 34 • Validation of the standard.
35 No proposed change and/or extension to the PMIx standard is accepted without an accompanying
36 prototype implementation in the PRI. This ensures that the proposal has undergone at least some
37 minimal level of scrutiny and testing before being considered.

- Ease of adoption.
The PRI is designed to be particularly easy for resource managers (and the SMS in general) to adopt, thus facilitating a rapid uptake into that community for application portability. Both client and server PMIx libraries are included, along with examples of client usage and server-side integration. A list of supported environments and versions is maintained on the PMIx web site <https://pmix.org/support/faq/what-apis-are-supported-on-my-rm/>

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

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

1.3.2 The PMIx Reference RunTime Environment (PRRTE)

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

1.4 Organization of this document

The remainder of this document is structured as follows:

- Introduction and Overview in Chapter 1 on page 1
- Terms and Conventions in Chapter 2 on page 11
- Data Structures and Types in Chapter 3 on page 16

- 1 • PMIx Initialization and Finalization in Chapter 4 on page 86
- 2 • Key/Value Management in Chapter 5 on page 99
- 3 • Process Management in Chapter 6 on page 129
- 4 • Job Management in Chapter 7 on page 148
- 5 • Event Notification in Chapter 8 on page 167
- 6 • Data Packing and Unpacking in Chapter 9 on page 176
- 7 • PMIx Server Specific Interfaces in Chapter 10 on page 186

8 1.5 Version 1.0: June 12, 2015

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

- 12 • Client APIs
 - 13 – PMIx_Init, **PMIx_Initialized**, **PMIx_Abort**, **PMIx_Finalize**
 - 14 – **PMIx_Put**, **PMIx_Commit**,
 - 15 – **PMIx_Fence**, **PMIx_Fence_nb**
 - 16 – **PMIx_Get**, **PMIx_Get_nb**
 - 17 – **PMIx_Publish**, **PMIx_Publish_nb**
 - 18 – **PMIx_Lookup**, **PMIx_Lookup**
 - 19 – **PMIx_Unpublish**, **PMIx_Unpublish_nb**
 - 20 – **PMIx_Spawn**, **PMIx_Spawn_nb**
 - 21 – **PMIx_Connect**, **PMIx_Connect_nb**
 - 22 – **PMIx_Disconnect**, **PMIx_Disconnect_nb**
 - 23 – **PMIx_Resolve_nodes**, **PMIx_Resolve_peers**
- 24 • Server APIs
 - 25 – **PMIx_server_init**, **PMIx_server_finalize**
 - 26 – **PMIx_generate_regex**, **PMIx_generate_ppn**
 - 27 – **PMIx_server_register_namespace**, **PMIx_server_deregister_namespace**
 - 28 – **PMIx_server_register_client**, **PMIx_server_deregister_client**

- 1 – `PMIx_server_setup_fork`, `PMIx_server_dmodex_request`
- 2 • Common APIs
- 3 – `PMIx_Get_version`, `PMIx_Store_internal`, `PMIx_Error_string`
- 4 – `PMIx_Register_errhandler`, `PMIx_Deregister_errhandler`,
- 5 `PMIx_Notify_error`

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

7 1.6 Version 2.0: Sept. 2018

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

- 9 • Client APIs
- 10 – `PMIx_Query_info_nb`, `PMIx_Log_nb`
- 11 – `PMIx_Allocation_request_nb`, `PMIx_Job_control_nb`,
- 12 `PMIx_Process_monitor_nb`, `PMIx_Heartbeat`
- 13 • Server APIs
- 14 – `PMIx_server_setup_application`, `PMIx_server_setup_local_support`
- 15 • Tool APIs
- 16 – `PMIx_tool_init`, `PMIx_tool_finalize`
- 17 • Common APIs
- 18 – `PMIx_Register_event_handler`, `PMIx_Deregister_event_handler`
- 19 – `PMIx_Notify_event`
- 20 – `PMIx_Proc_state_string`, `PMIx_Scope_string`
- 21 – `PMIx_Persistence_string`, `PMIx_Data_range_string`
- 22 – `PMIx_Info_directives_string`, `PMIx_Data_type_string`
- 23 – `PMIx_Alloc_directive_string`
- 24 – `PMIx_Data_pack`, `PMIx_Data_unpack`, `PMIx_Data_copy`
- 25 – `PMIx_Data_print`, `PMIx_Data_copy_payload`

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

1 1.7 Version 2.1: Dec. 2018

2 The v2.1 update includes clarifications and corrections, plus addition of examples:

- 3 • Clarify description of `PMIx_Connect` and `PMIx_Disconnect` APIs.
- 4 • Explain that values for the `PMIX_COLLECTIVE_ALGO` are environment-dependent
- 5 • Identify the namespace/rank values required for retrieving attribute-associated information using
6 the `PMIx_Get` API
- 7 • Provide definitions for `session` , `job` , `application` , and other terms used throughout the
8 document
- 9 • Clarify definitions of `PMIX_UNIV_SIZE` versus `PMIX_JOB_SIZE`
- 10 • Clarify server module function return values
- 11 • Provide examples of the use of `PMIx_Get` for retrieval of information
- 12 • Clarify the use of `PMIx_Get` versus `PMIx_Query_info_nb`
- 13 • Clarify return values for non-blocking APIs and emphasize that callback functions must not be
14 invoked prior to return from the API
- 15 • Provide detailed example for construction of the `PMIx_server_register_namespace` input
16 information array
- 17 • Define information levels (e.g., `session` vs `job`) and associated attributes for both storing
18 and retrieving values
- 19 • Clarify roles of PMIx server library and host environment for collective operations
- 20 • Clarify definition of `PMIX_UNIV_SIZE`

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, “*mpirexec -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 • *resource manager* is used in a generic sense to represent the system that will host the PMIx
2 server library. This could be a vendor's RM, a programming library's RunTime
3 Environment (RTE), or some other agent.
- 4 • *host environment* is used interchangeably with *resource manager* to refer to the process hosting
5 the PMIx server library.

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

10 2.1 Notational Conventions

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

▼  ▼

13 C specific text...

14 `int foo = 42;`

▲  ▲

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

▼  ▼

17 Note: General text...

▲  ▲

▼  ▼

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

▲  ▲

▼  ▼

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

▲  ▲

Advice to PMIx library implementers

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

Advice to PMIx server hosts

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

2.2 Semantics

The following terms will be taken to mean:

- *shall* and *will* indicate that the specified behavior is *required* of all conforming implementations
- *should* and *may* indicate behaviors that a quality implementation would include, but are not required of all conforming implementations

2.3 Naming Conventions

The PMIx standard has adopted the following conventions:

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

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

1 2.4 Procedure Conventions

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

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

11 2.5 Standard vs Reference Implementation

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

18 Similarly, the *PMIx Reference RunTime Environment* (PRRTE) is provided by the community to
19 enable users operating in non-PMIx environments to develop and execute PMIx-enabled
20 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, types, and constants. These apply to all
2 consumers of the PMIx interface. Where necessary for clarification, the description of, for
3 example, an attribute may be copied from this chapter into a section where it is used.

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

▼ Advice to PMIx library implementers ▼

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

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

12 3.1 Constants

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

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

▼ Advice to PMIx library implementers ▼

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

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

Advice to PMIx library implementers

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

4 3.1.1 Error Constants

5 The **pmix_status_t** structure is an **int** type for return status.

6 The tables shown in this section define the possible values for **pmix_status_t** . PMIx errors are
7 required to always be negative, with 0 reserved for **PMIX_SUCCESS** .

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

Advice to users

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

1 3.1.1.1 PMIx v1 Error Constants

2 The following list contains those constants defined in the PMIx v1 standard. Those values in the list
3 that were deprecated in later standards are denoted as such. PMIx errors are always negative, with 0
4 reserved for success.

5	PMIX_SUCCESS	Success
6	PMIX_ERROR	General Error
7	PMIX_ERR_SILENT	Silent error
8	PMIX_ERR_DEBUGGER_RELEASE	Error in debugger release
9	PMIX_ERR_PROC_RESTART	Fault tolerance: Error in process restart
10	PMIX_ERR_PROC_CHECKPOINT	Fault tolerance: Error in process checkpoint
11	PMIX_ERR_PROC_MIGRATE	Fault tolerance: Error in process migration
12	PMIX_ERR_PROC_ABORTED	Process was aborted
13	PMIX_ERR_PROC_REQUESTED_ABORT	Process is already requested to abort
14	PMIX_ERR_PROC_ABORTING	Process is being aborted
15	PMIX_ERR_SERVER_FAILED_REQUEST	Failed to connect to the server
16	PMIX_EXISTS	Requested operation would overwrite an existing value
17	PMIX_ERR_INVALID_CRED	Invalid security credentials
18	PMIX_ERR_HANDSHAKE_FAILED	Connection handshake failed
19	PMIX_ERR_READY_FOR_HANDSHAKE	Ready for handshake
20	PMIX_ERR_WOULD_BLOCK	Operation would block
21	PMIX_ERR_UNKNOWN_DATA_TYPE	Unknown data type
22	PMIX_ERR_PROC_ENTRY_NOT_FOUND	Process not found
23	PMIX_ERR_TYPE_MISMATCH	Invalid type
24	PMIX_ERR_UNPACK_INADEQUATE_SPACE	Inadequate space to unpack data
25	PMIX_ERR_UNPACK_FAILURE	Unpack failed
26	PMIX_ERR_PACK_FAILURE	Pack failed
27	PMIX_ERR_PACK_MISMATCH	Pack mismatch
28	PMIX_ERR_NO_PERMISSIONS	No permissions
29	PMIX_ERR_TIMEOUT	Timeout expired
30	PMIX_ERR_UNREACH	Unreachable
31	PMIX_ERR_IN_ERRNO	Error defined in errno
32	PMIX_ERR_BAD_PARAM	Bad parameter
33	PMIX_ERR_RESOURCE_BUSY	Resource busy
34	PMIX_ERR_OUT_OF_RESOURCE	Resource exhausted
35	PMIX_ERR_DATA_VALUE_NOT_FOUND	Data value not found
36	PMIX_ERR_INIT	Error during initialization
37	PMIX_ERR_NOMEM	Out of memory
38	PMIX_ERR_INVALID_ARG	Invalid argument
39	PMIX_ERR_INVALID_KEY	Invalid key
40	PMIX_ERR_INVALID_KEY_LENGTH	Invalid key length
41	PMIX_ERR_INVALID_VAL	Invalid value

1 **PMIX_ERR_INVALID_VAL_LENGTH** Invalid value length
 2 **PMIX_ERR_INVALID_LENGTH** Invalid argument length
 3 **PMIX_ERR_INVALID_NUM_ARGS** Invalid number of arguments
 4 **PMIX_ERR_INVALID_ARGS** Invalid arguments
 5 **PMIX_ERR_INVALID_NUM_PARSED** Invalid number parsed
 6 **PMIX_ERR_INVALID_KEYVALP** Invalid key/value pair
 7 **PMIX_ERR_INVALID_SIZE** Invalid size
 8 **PMIX_ERR_INVALID_NAMESPACE** Invalid namespace
 9 **PMIX_ERR_SERVER_NOT_AVAIL** Server is not available
 10 **PMIX_ERR_NOT_FOUND** Not found
 11 **PMIX_ERR_NOT_SUPPORTED** Not supported
 12 **PMIX_ERR_NOT_IMPLEMENTED** Not implemented
 13 **PMIX_ERR_COMM_FAILURE** Communication failure
 14 **PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER** Unpacking past the end of the buffer
 15 provided

16 3.1.1.2 PMIx v2 Error Constants

17 The following list contains constants added in the PMIx v2 standard.

18 **PMIX_ERR_LOST_CONNECTION_TO_SERVER** Lost connection to server
 19 **PMIX_ERR_LOST_PEER_CONNECTION** Lost connection to peer
 20 **PMIX_ERR_LOST_CONNECTION_TO_CLIENT** Lost connection to client
 21 **PMIX_QUERY_PARTIAL_SUCCESS** Query partial success (used by query system)
 22 **PMIX_NOTIFY_ALLOC_COMPLETE** Notify that allocation is complete
 23 **PMIX_JCTRL_CHECKPOINT** Job control: Monitored by PMIx client to trigger checkpoint
 24 operation
 25 **PMIX_JCTRL_CHECKPOINT_COMPLETE** Job control: Sent by PMIx client and monitored
 26 by PMIx server to notify that requested checkpoint operation has completed.
 27 **PMIX_JCTRL_PREEMPT_ALERT** Job control: Monitored by PMIx client to detect an RM
 28 intending to preempt the job.
 29 **PMIX_MONITOR_HEARTBEAT_ALERT** Job monitoring: Heartbeat alert
 30 **PMIX_MONITOR_FILE_ALERT** Job monitoring: File alert
 31 **PMIX_PROC_TERMINATED** Process terminated - can be either normal or abnormal
 32 termination
 33 **PMIX_ERR_INVALID_TERMINATION** Process terminated without calling
 34 **PMIx_Finalize**, or was a member of an assemblage formed via **PMIx_Connect** and
 35 terminated or called **PMIx_Finalize** without first calling **PMIx_Disconnect** (or its
 36 non-blocking form) from that assemblage.

37 The following list contains operational error constants introduced in the v2 standard.

38 **PMIX_ERR_EVENT_REGISTRATION** Error in event registration
 39 **PMIX_ERR_JOB_TERMINATED** Error job terminated
 40 **PMIX_ERR_UPDATE_ENDPOINTS** Error updating endpoints

1 **PMIX_MODEL_DECLARED** Model declared
2 **PMIX_GDS_ACTION_COMPLETE** The global data storage (GDS) action has completed
3 **PMIX_ERR_INVALID_OPERATION** The requested operation is supported by the
4 implementation and host environment, but fails to meet a requirement (e.g., requesting to
5 *disconnect* from processes without first *connecting* to them).

6 The following list contains system error constants introduced in the v2 standard.

7 **PMIX_ERR_NODE_DOWN** Node down
8 **PMIX_ERR_NODE_OFFLINE** Node is marked as offline

9 The following list contains event handler error constants introduced in the v2 standard.

10 **PMIX_EVENT_NO_ACTION_TAKEN** Event handler: No action taken
11 **PMIX_EVENT_PARTIAL_ACTION_TAKEN** Event handler: Partial action taken
12 **PMIX_EVENT_ACTION_DEFERRED** Event handler: Action deferred
13 **PMIX_EVENT_ACTION_COMPLETE** Event handler: Action complete

14 3.1.1.3 User-Defined Error Constants

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

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

22 3.2 Data Types

23 This section defines various data types used by the PMIx APIs.

24 3.2.1 Key Structure

25 The **pmix_key_t** structure is a statically defined character array of length **PMIX_MAX_KEYLEN**
26 +1, thus supporting keys of maximum length **PMIX_MAX_KEYLEN** while preserving space for a
27 mandatory **NULL** terminator.

PMIx v2.0

 C

28 **typedef char pmix_key_t[PMIX_MAX_KEYLEN+1];**

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

Advice to users

3 References to keys in PMIx v1 rwere defined simply as an array of characters of size
 4 **PMIX_MAX_KEYLEN+1**. The **pmix_key_t** type definition was introduced in version 2 of the
 5 standard. The two definitions are code-compatible and thus do not represent a break in backward
 6 compatibility.

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

10 3.2.2 Namespace Structure

11 The **pmix_namespace_t** structure is a statically defined character array of length
 12 **PMIX_MAX_NSLEN+1**, thus supporting namespaces of maximum length **PMIX_MAX_NSLEN**
 13 while preserving space for a mandatory **NULL** terminator.

PMIx v2.0

```
14 typedef char pmix_namespace_t [PMIX_MAX_NSLEN+1];
```

15 Characters in the namespace must be standard alphanumeric values supported by common utilities
 16 such as *strcmp*.

Advice to users

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

21 Passing a **pmix_namespace_t** value to the standard *sizeof* utility can result in compiler warnings of
 22 incorrect returned value. Users are advised to avoid using *sizeof(pmix_namespace_t)* and instead rely
 23 on the **PMIX_MAX_NSLEN** constant.

1 3.2.3 Rank Structure

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

PMIx v1.0

```
▼ _____ C _____ ▼  
3 typedef uint32_t pmix_rank_t;  
▲ _____ C _____ ▲
```

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 v2.0* **PMIX_RANK_LOCAL_NODE** Special rank value used to define groups of ranks for use in
13 collectives. This constant defines the group of all ranks on a local node.

14 3.2.4 Process Structure

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

PMIx v1.0

```
▼ _____ C _____ ▼  
17 typedef struct pmix_proc {  
18     pmix_namespace_t nspace;  
19     pmix_rank_t rank;  
20 } pmix_proc_t;  
▲ _____ C _____ ▲
```

21 3.2.5 Process structure support macros

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

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

2 Initialize the `pmix_proc_t` fields

PMIx v1.0



C

3 **PMIX_PROC_CONSTRUCT (m)**



C

4 **IN m**

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

6 **3.2.5.2 Destruct the `pmix_proc_t` structure**

7 Clear the `pmix_proc_t` fields

PMIx v1.0



C

8 **PMIX_PROC_DESTRUCT (m)**



C

9 **IN m**

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

11 This macro performs the identical operations as `PMIX_PROC_CONSTRUCT`, but is provided for
12 symmetry in user code.

13 **3.2.5.3 Create a `pmix_proc_t` array**

14 Allocate and initialize an array of `pmix_proc_t` structures

PMIx v1.0



C

15 **PMIX_PROC_CREATE (m, n)**



C

16 **INOUT m**

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

18 **IN n**

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

1 3.2.5.4 Free a `pmix_proc_t` array

2 Release an array of `pmix_proc_t` structures

PMIx v1.0

▼  C 

3 **PMIX_PROC_FREE** (m, n)

▲  C 

4 **IN** m

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

6 **IN** n

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

8 3.2.5.5 Load a `pmix_proc_t` structure

9 Load values into a `pmix_proc_t`

PMIx v2.0

▼  C 

10 **PMIX_PROC_LOAD** (m, n, r)

▲  C 

11 **IN** m

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

13 **IN** n

14 Namespace to be loaded (`pmix_namespace_t`)

15 **IN** r

16 Rank to be assigned (`pmix_rank_t`)

17 3.2.6 Process State Structure

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

▼  Advice to users 

21 The fine-grained nature of the following constants may exceed the ability of an RM to provide
22 updated process state values during the process lifetime. This is particularly true of states in the
23 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 this constant and
9 **PMIX_PROC_STATE_CONNECTED** so users can easily and quickly determine if a process
10 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;

```

15 3.2.8 Process Information Structure support macros

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

17 3.2.8.1 Initialize the `pmix_proc_info_t` structure

18 Initialize the `pmix_proc_info_t` fields

PMIx v2.0

```

19 #define PMIX_PROC_INFO_CONSTRUCT(m)

```

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

22 3.2.8.2 Destruct the `pmix_proc_info_t` structure

23 Destruct the `pmix_proc_info_t` fields

PMIx v2.0

```

24 #define PMIX_PROC_INFO_DESTRUCT(m)

```

25 **IN** `m`
26 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

▼ `PMIX_PROC_INFO_CREATE` C ▼

3 `PMIX_PROC_INFO_CREATE` (*m*, *n*)

▲ `PMIX_PROC_INFO_CREATE` 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

▼ `PMIX_PROC_INFO_FREE` C ▼

11 `PMIX_PROC_INFO_FREE` (*m*, *n*)

▲ `PMIX_PROC_INFO_FREE` 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.

1 **PMIX_GLOBAL** The data is to be shared with all other requesting processes, regardless of
 2 location.
 3 *PMIx v2.0* **PMIX_INTERNAL** The data is intended solely for this process and is not shared with other
 4 processes.

5 **3.2.10 Range of Published Data**

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

11 **PMIX_RANGE_UNDEF** Undefined range
 12 **PMIX_RANGE_RM** Data is intended for the host resource manager.
 13 **PMIX_RANGE_LOCAL** Data is only available to processes on the local node.
 14 **PMIX_RANGE_NAMESPACE** Data is only available to processes in the same namespace.
 15 **PMIX_RANGE_SESSION** Data is only available to all processes in the session.
 16 **PMIX_RANGE_GLOBAL** Data is available to all processes.
 17 **PMIX_RANGE_CUSTOM** Range is specified in the **pmix_info_t** associated with this call.
 18 **PMIX_RANGE_PROC_LOCAL** Data is only available to this process.

▼ Advice to users ▼

19 The names of the **pmix_data_range_t** values changed between version 1 and version 2 of the
 20 standard, thereby breaking backward compatibility

21 **3.2.11 Data Persistence Structure**

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

26 **PMIX_PERSIST_INDEF** Retain data until specifically deleted.
 27 **PMIX_PERSIST_FIRST_READ** Retain data until the first access, then the data is deleted.
 28 **PMIX_PERSIST_PROC** Retain data until the publishing process terminates.
 29 **PMIX_PERSIST_APP** Retain data until the application terminates.
 30 **PMIX_PERSIST_SESSION** Retain data until the session/allocation terminates.

1 3.2.12 Value Structure

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

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

PMIx v1.0

C

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

```

1         pmix_alloc_directive_t adir;    // version 2.0
2         /**** DEPRECATED in PMIx 2 ****/
3         pmix_info_array_t *array;
4         /*****
5         } data;
6     } pmix_value_t;

```



7 3.2.13 Value structure support macros

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

9 3.2.13.1 Initialize the `pmix_value_t` structure

10 Initialize the `pmix_value_t` fields

PMIx v1.0



11 **PMIX_VALUE_CONSTRUCT** (m)



12 **IN** m

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

14 3.2.13.2 Destruct the `pmix_value_t` structure

15 Destruct the `pmix_value_t` fields

PMIx v1.0



16 **PMIX_VALUE_DESTRUCT** (m)



17 **IN** m

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

19 3.2.13.3 Create a `pmix_value_t` array

20 Allocate and initialize an array of `pmix_value_t` structures

PMIx v1.0



21 **PMIX_VALUE_CREATE** (m, n)



22 **INOUT** m

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

25 **IN** n

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

1 3.2.13.4 Free a `pmix_value_t` array

2 Release an array of `pmix_value_t` structures

PMIx v1.0

▼ `PMIX_VALUE_FREE` C

3 `PMIX_VALUE_FREE(m, n)`

▲ `PMIX_VALUE_FREE` C

4 **IN** `m`

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

6 **IN** `n`

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

8 3.2.14 Load a `pmix_value_t` structure

9 Summary

10 Load data into a `pmix_value_t` structure.

PMIx v2.0

▼ `PMIX_VALUE_LOAD` C

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

▲ `PMIX_VALUE_LOAD` C

12 **IN** `v`

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

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 Description

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

▼ **Advice to users** ▼

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

▲

1 3.2.14.1 Transfer data between `pmix_value_t` structures

2 Summary

3 Transfer the data value between two `pmix_value_t` structures.

PMIx v2.0

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

5 **OUT** r

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

7 **IN** d

8 Pointer to the `pmix_value_t` destination (handle)

9 **IN** s

10 Pointer to the `pmix_value_t` source (handle)

11 Description

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

Advice to users

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

16 3.2.15 Info and Info Array Structures

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

PMIx v1.0

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

24 The `pmix_info_array` structure defines an array of `pmix_info_t` structures.

1 Note: The `pmix_info_array` structure has been deprecated and will be removed in future
2 versions of the PMIx Standard.

PMIx v1.0

```
3 typedef struct pmix_info_array {  
4     size_t size;  
5     pmix_info_t *array;  
6 } pmix_info_array_t;
```

7 3.2.16 Info structure support macros

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

9 3.2.16.1 Initialize the `pmix_info_t` structure

10 Initialize the `pmix_info_t` fields

PMIx v1.0

```
11 PMIX_INFO_CONSTRUCT (m)
```

12 **IN** m

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

14 3.2.16.2 Destruct the `pmix_info_t` structure

15 Destruct the `pmix_info_t` fields

PMIx v1.0

```
16 PMIX_INFO_DESTRUCT (m)
```

17 **IN** m

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

1 3.2.16.3 Create a `pmix_info_t` array

2 Allocate and initialize an array of `pmix_info_t` structures

PMIx v1.0

▼ `PMIX_INFO_CREATE` C `PMIX_INFO_CREATE` ▼

3 `PMIX_INFO_CREATE(m, n)`

▲ `PMIX_INFO_CREATE` C `PMIX_INFO_CREATE` ▲

4 **INOUT** `m`

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

6 **IN** `n`

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

8 3.2.16.4 Free a `pmix_info_t` array

9 Release an array of `pmix_info_t` structures

PMIx v1.0

▼ `PMIX_INFO_FREE` C `PMIX_INFO_FREE` ▼

10 `PMIX_INFO_FREE(m, n)`

▲ `PMIX_INFO_FREE` C `PMIX_INFO_FREE` ▲

11 **IN** `m`

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

13 **IN** `n`

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

15 3.2.16.5 Load key and value data into a `pmix_info_t`

PMIx v1.0

▼ `PMIX_INFO_LOAD` C `PMIX_INFO_LOAD` ▼

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

▲ `PMIX_INFO_LOAD` C `PMIX_INFO_LOAD` ▲

17 **IN** `v`

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

20 **IN** `k`

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

23 **IN** `d`

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

25 **IN** `t`

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

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

Advice to users

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

3.2.16.6 Copy data between `pmix_info_t` structures

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

PMIx v2.0

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

IN `d`

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

IN `s`

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

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

Advice to users

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

3.2.16.7 Test a boolean `pmix_info_t`

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

PMIx v2.0

```
PMIX_INFO_TRUE(m)
```

IN `m`

Pointer to a `pmix_info_t` structure (handle)

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

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

1 3.2.17 Info Type Directives

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

▼ Advice to users ▼

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

▼ Advice to PMIx library implementers ▼

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

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

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

▼ Advice to PMIx server hosts ▼

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

23 3.2.18 Info Directive support macros

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

1 3.2.18.1 Mark an info structure as required

2 Summary

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

PMIx v2.0

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

5 **IN** `info`

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

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

8 3.2.18.2 Test an info structure for *required* directive

9 Summary

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

PMIx v2.0

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

12 **IN** `info`

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

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

15 3.2.19 Job Allocation Directives

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

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

22 **PMIX_ALLOC_EXTEND** Extend the existing allocation, either in time or as additional
23 resources.

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

27 **PMIX_ALLOC_REAQUIRE** Reacquire resources that were previously “lent” back to the
28 scheduler.

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

1 3.2.20 Lookup Returned Data Structure

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

PMIx v1.0

```
3 typedef struct pmix_pdata {  
4     pmix_proc_t proc;  
5     pmix_key_t key;  
6     pmix_value_t value;  
7 } pmix_pdata_t;
```

8 3.2.21 Lookup data structure support macros

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

10 3.2.21.1 Initialize the `pmix_pdata_t` structure

11 Initialize the `pmix_pdata_t` fields

PMIx v1.0

```
12 PMIX_PDATA_CONSTRUCT (m)
```

13 **IN** m

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

15 3.2.21.2 Destruct the `pmix_pdata_t` structure

16 Destruct the `pmix_pdata_t` fields

PMIx v1.0

```
17 PMIX_PDATA_DESTRUCT (m)
```

18 **IN** m

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

1 3.2.21.3 Create a `pmix_pdata_t` array

2 Allocate and initialize an array of `pmix_pdata_t` structures

PMIx v1.0

▼ `PMIX_PDATA_CREATE` C ▲

3 `PMIX_PDATA_CREATE` (*m*, *n*)

▲ `PMIX_PDATA_CREATE` C ▲

4 **INOUT** *m*

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

7 **IN** *n*

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

9 3.2.21.4 Free a `pmix_pdata_t` array

10 Release an array of `pmix_pdata_t` structures

PMIx v1.0

▼ `PMIX_PDATA_FREE` C ▲

11 `PMIX_PDATA_FREE` (*m*, *n*)

▲ `PMIX_PDATA_FREE` C ▲

12 **IN** *m*

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

14 **IN** *n*

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

16 3.2.21.5 Load a lookup data structure

17 Summary

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

PMIx v1.0

▼ `PMIX_PDATA_LOAD` C ▲

19 `PMIX_PDATA_LOAD` (*m*, *p*, *k*, *d*, *t*);

▲ `PMIX_PDATA_LOAD` C ▲

20 **IN** *m*

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

23 **IN** *p*

24 Pointer to the `pmix_proc_t` structure containing the identifier of the process being
25 referenced (pointer to `pmix_proc_t`)

26 **IN** *k*

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

1 **IN d**
2 Pointer to the data value to be loaded (handle)

3 **IN t**
4 Type of the provided data value (`pmix_data_type_t`)

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

▼ **Advice to users** ▼

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



10 3.2.21.6 Transfer a lookup data structure

11 Summary

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

PMIx v2.0



C

13 **PMIX_PDATA_XFER(d, s);**



C

14 **IN d**
15 Pointer to the destination `pmix_pdata_t` (pointer to `pmix_pdata_t`)

16 **IN s**
17 Pointer to the source `pmix_pdata_t` (pointer to `pmix_pdata_t`)

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

▼ **Advice to users** ▼

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



1 3.2.22 Application Structure

2 The `pmix_app_t` structure describes the application context for the `PMIx_Spawn` and
3 `PMIx_Spawn_nb` operations.

PMIx v1.0

C

```
4 typedef struct pmix_app {  
5     /** Executable */  
6     char *cmd;  
7     /** Argument set, NULL terminated */  
8     char **argv;  
9     /** Environment set, NULL terminated */  
10    char **env;  
11    /** Current working directory */  
12    char *cwd;  
13    /** Maximum processes with this profile */  
14    int maxprocs;  
15    /** Array of info keys describing this application*/  
16    pmix_info_t *info;  
17    /** Number of info keys in 'info' array */  
18    size_t ninfo;  
19 } pmix_app_t;
```

C

20 3.2.23 App structure support macros

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

22 3.2.23.1 Initialize the `pmix_app_t` structure

23 Initialize the `pmix_app_t` fields

PMIx v1.0

C

```
24 PMIX_APP_CONSTRUCT(m)
```

C

25 **IN** m

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

1 **3.2.23.2 Destruct the `pmix_app_t` structure**

2 Destruct the `pmix_app_t` fields

PMIx v1.0



3 **PMIX_APP_DESTRUCT (m)**



4 **IN m**

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

6 **3.2.23.3 Create a `pmix_app_t` array**

7 Allocate and initialize an array of `pmix_app_t` structures

PMIx v1.0



8 **PMIX_APP_CREATE (m, n)**



9 **INOUT m**

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

11 **IN n**

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

13 **3.2.23.4 Free a `pmix_app_t` array**

14 Release an array of `pmix_app_t` structures

PMIx v1.0



15 **PMIX_APP_FREE (m, n)**



16 **IN m**

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

18 **IN n**

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

1 3.2.24 Query Structure

2 The `pmix_query_t` structure is used by `PMIx_Query_info_nb` to describe a single query
3 operation.

PMIx v2.0

```
4 typedef struct pmix_query {  
5     char **keys;  
6     pmix_info_t *qualifiers;  
7     size_t nqual;  
8 } pmix_query_t;
```

C

C

9 3.2.25 Query structure support macros

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

11 3.2.25.1 Initialize the `pmix_query_t` structure

12 Initialize the `pmix_query_t` fields

PMIx v2.0

```
13 PMIX_QUERY_CONSTRUCT (m)
```

14 **IN** m

15 Pointer to the structure to be initialized (pointer to `pmix_query_t`)

16 3.2.25.2 Destruct the `pmix_query_t` structure

17 Destruct the `pmix_query_t` fields

PMIx v2.0

```
18 PMIX_QUERY_DESTRUCT (m)
```

19 **IN** m

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

1 3.2.25.3 Create a `pmix_query_t` array

2 Allocate and initialize an array of `pmix_query_t` structures

PMIx v2.0

▼  ▼

3 `PMIX_QUERY_CREATE(m, n)`

▲  ▲

4 **INOUT** `m`

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

7 **IN** `n`

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

9 3.2.25.4 Free a `pmix_query_t` array

10 Release an array of `pmix_query_t` structures

PMIx v2.0

▼  ▼

11 `PMIX_QUERY_FREE(m, n)`

▲  ▲

12 **IN** `m`

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

14 **IN** `n`

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

16 3.2.26 Modex Structure

17 The `pmix_modex_data_t` structure describes the business card exchange (BCX) information.

▼ 

18 Note: This structure and its supporting macros have been deprecated and will be removed in future
19 versions of the PMIx Standard.

▲ 

PMIx v1.0

▼  ▼

```
20 typedef struct pmix_modex_data {  
21     pmix_nspace_t nspace;  
22     int rank;  
23     uint8_t *blob;  
24     size_t size;  
25 } pmix_modex_data_t;
```

▲  ▲

1 3.2.27 Modex data structure support macros

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

3 3.2.27.1 Initialize the `pmix_modex_data_t` structure

4 Initialize the `pmix_modex_data_t` fields

PMIx v1.0

▼  

5 **PMIX_MODEX_CONSTRUCT** (m)

▲  

6 **IN** m

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

8 3.2.27.2 Destruct the `pmix_modex_data_t` structure

9 Destruct the `pmix_modex_data_t` fields

PMIx v1.0

▼  

10 **PMIX_MODEX_DESTRUCT** (m)

▲  

11 **IN** m

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

13 3.2.27.3 Create a `pmix_modex_data_t` array

14 Allocate and initialize an array of `pmix_modex_data_t` structures

PMIx v1.0

▼  

15 **PMIX_MODEX_CREATE** (m, n)

▲  

16 **INOUT** m

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

19 **IN** n

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

1 3.2.27.4 Free a `pmix_modex_data_t` array

2 Release an array of `pmix_modex_data_t` structures

PMIx v1.0

▼  ▼

3 `PMIX_MODEX_FREE(m, n)`

▲  ▲

4 **IN** `m`

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

6 **IN** `n`

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

8 3.3 Data Packing/Unpacking Types and Structures

9 This section defines types and structures used to pack and unpack data passed through the PMIx
10 API.

11 3.3.1 Byte Object Type

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

PMIx v1.0

▼  ▼

```
13 typedef struct pmix_byte_object {  
14     char *bytes;  
15     size_t size;  
16 } pmix_byte_object_t;
```

▲  ▲

17 3.3.2 Byte object support macros

18 The following macros support the `pmix_byte_object_t` structure.

19 3.3.2.1 Initialize the `pmix_byte_object_t` structure

20 Initialize the `pmix_byte_object_t` fields

PMIx v2.0

▼  ▼

21 `PMIX_BYTE_OBJECT_CONSTRUCT(m)`

▲  ▲

22 **IN** `m`

23 Pointer to the structure to be initialized (pointer to `pmix_byte_object_t`)

1 3.3.2.2 Destruct the `pmix_byte_object_t` structure

2 Clear the `pmix_byte_object_t` fields

PMIx v2.0

▼ `PMIX_BYTE_OBJECT_DESTRUCT` (m) C

3 `PMIX_BYTE_OBJECT_DESTRUCT` (m)

▲ C

4 **IN** m

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

6 3.3.2.3 Create a `pmix_byte_object_t` structure

7 Allocate and initialize an array of `pmix_byte_object_t` structures

PMIx v2.0

▼ `PMIX_BYTE_OBJECT_CREATE` (m, n) C

8 `PMIX_BYTE_OBJECT_CREATE` (m, n)

▲ C

9 **INOUT** m

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

12 **IN** n

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

14 3.3.2.4 Free a `pmix_byte_object_t` array

15 Release an array of `pmix_byte_object_t` structures

PMIx v2.0

▼ `PMIX_BYTE_OBJECT_FREE` (m, n) C

16 `PMIX_BYTE_OBJECT_FREE` (m, n)

▲ C

17 **IN** m

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

19 **IN** n

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

1 3.3.2.5 Load a `pmix_byte_object_t` structure

2 Load values into a `pmix_byte_object_t`

PMIx v2.0

▼  

3 `PMIX_BYTE_OBJECT_LOAD(b, d, s)`

▲  

4 **IN** `b`

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

6 **IN** `d`

7 Pointer to the data to be loaded (`char*`)

8 **IN** `s`

9 Number of bytes in the data array (`size_t`)

10 3.3.3 Data Buffer Type

11 The `pmix_data_buffer_t` structure describes a data buffer used for packing and unpacking.

PMIx v2.0

▼  

```
12 typedef struct pmix_data_buffer {
13     /** Start of my memory */
14     char *base_ptr;
15     /** Where the next data will be packed to (within the allocated
16         memory starting at base_ptr) */
17     char *pack_ptr;
18     /** Where the next data will be unpacked from (within the
19         allocated memory starting as base_ptr) */
20     char *unpack_ptr;
21     /** Number of bytes allocated (starting at base_ptr) */
22     size_t bytes_allocated;
23     /** Number of bytes used by the buffer (i.e., amount of data --
24         including overhead -- packed in the buffer) */
25     size_t bytes_used;
26 } pmix_data_buffer_t;
```

▲  

27 3.3.4 Data buffer support macros

28 The following macros support the `pmix_data_buffer_t` structure.

1 3.3.4.1 Initialize the `pmix_data_buffer_t` structure

2 Initialize the `pmix_data_buffer_t` fields

PMIx v2.0

▼  

3 `PMIX_DATA_BUFFER_CONSTRUCT (m)`

▲  

4 **IN** `m`

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

6 3.3.4.2 Destruct the `pmix_data_buffer_t` structure

7 Clear the `pmix_data_buffer_t` fields

PMIx v2.0

▼  

8 `PMIX_DATA_BUFFER_DESTRUCT (m)`

▲  

9 **IN** `m`

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

11 3.3.4.3 Create a `pmix_data_buffer_t` structure

12 Allocate and initialize a `pmix_data_buffer_t` structure

PMIx v2.0

▼  

13 `PMIX_DATA_BUFFER_CREATE (m)`

▲  

14 **INOUT** `m`

15 Address where the pointer to the `pmix_data_buffer_t` structure shall be stored
16 (handle)

17 3.3.4.4 Free a `pmix_data_buffer_t`

18 Release a `pmix_data_buffer_t` structure

PMIx v2.0

▼  

19 `PMIX_DATA_BUFFER_RELEASE (m)`

▲  

20 **IN** `m`

21 Pointer to the `pmix_data_buffer_t` structure to be released (handle)

1 3.3.5 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;
```

C

C

8 3.3.6 Generalized Data Types Used for Packing/Unpacking

9 The `pmix_data_type_t` structure is a `uint16_t` type for identifying the data type for
10 packing/unpacking purposes.

Advice to PMIx library implementers

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

16 3.3.6.1 PMIx v1 Data Types

17 The following types were introduced in version 1 of the PMIx Standard.

18	<code>PMIX_UNDEF</code>	Undefined
19	<code>PMIX_BOOL</code>	Boolean (converted to/from native <code>true/false</code>) (<code>bool</code>)
20	<code>PMIX_BYTE</code>	A byte of data (<code>uint8_t</code>)
21	<code>PMIX_STRING</code>	<code>NULL</code> terminated string (<code>char*</code>)
22	<code>PMIX_SIZE</code>	Size <code>size_t</code>
23	<code>PMIX_PID</code>	Operating process identifier (PID) (<code>pid_t</code>)
24	<code>PMIX_INT</code>	Integer (<code>int</code>)
25	<code>PMIX_INT8</code>	8-byte integer (<code>int8_t</code>)
26	<code>PMIX_INT16</code>	16-byte integer (<code>int16_t</code>)
27	<code>PMIX_INT32</code>	32-byte integer (<code>int32_t</code>)
28	<code>PMIX_INT64</code>	64-byte integer (<code>int64_t</code>)
29	<code>PMIX_UINT</code>	Unsigned integer (<code>unsigned int</code>)
30	<code>PMIX_UINT8</code>	Unsigned 8-byte integer (<code>uint8_t</code>)

1 **PMIX_UINT16** Unsigned 16-byte integer (`uint16_t`)
 2 **PMIX_UINT32** Unsigned 32-byte integer (`uint32_t`)
 3 **PMIX_UINT64** Unsigned 64-byte integer (`uint64_t`)
 4 **PMIX_FLOAT** Float (`float`)
 5 **PMIX_DOUBLE** Double (`double`)
 6 **PMIX_TIMEVAL** Time value (`struct timeval`)
 7 **PMIX_TIME** Time (`time_t`)
 8 **PMIX_VALUE** Value (`pmix_value_t`)
 9 **PMIX_PROC** Process (`pmix_proc_t`)
 10 **PMIX_APP** Application context
 11 **PMIX_INFO** Info object
 12 **PMIX_PDATA** Pointer to data
 13 **PMIX_BUFFER** Buffer
 14 **PMIX_BYTE_OBJECT** Byte object (`pmix_byte_object_t`)
 15 **PMIX_KVAL** Key/value pair
 16 **PMIX_MODEX (Deprecated in PMIx 2.0)** Modex
 17 **PMIX_PERSIST** Persistence (`pmix_persistence_t`)
 18 **PMIX_INFO_ARRAY (Deprecated in PMIx 2.0)** Info array

19 3.3.6.2 PMIx v2 Data Types

20 The following types were introduced in version 2 of the PMIx Standard.

21 **PMIX_STATUS** Status (`pmix_status_t`)
 22 **PMIX_POINTER** Pointer (`void*`)
 23 **PMIX_SCOPE** Scope (`pmix_scope_t`)
 24 **PMIX_DATA_RANGE** Data range (`pmix_data_range_t`)
 25 **PMIX_COMMAND** Command
 26 **PMIX_INFO_DIRECTIVES** Info directives
 27 **PMIX_DATA_TYPE** Data type
 28 **PMIX_PROC_STATE** Process state (`pmix_proc_state_t`)
 29 **PMIX_PROC_INFO** Process info (`pmix_proc_info_t`)
 30 **PMIX_DATA_ARRAY** Data array (`pmix_data_array_t`)
 31 **PMIX_PROC_RANK** Process rank (`pmix_rank_t`)
 32 **PMIX_QUERY** Query
 33 **PMIX_COMPRESSED_STRING** Compressed string (with zlib)
 34 **PMIX_ALLOC_DIRECTIVE** Allocation directive (`pmix_alloc_directive_t`)
 35 **PMIX_DATA_TYPE_MAX** A boundary for implementers above which they can add their own
 36 data types.

1 3.4 Reserved attributes

2 The PMIx standard defines a relatively small set of APIs and the caller may customize the behavior
3 of the API by passing one or more attributes to that API. Additionally, attributes may be keys
4 passed to `PMIx_Get` calls to access the specified values from the system.

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

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

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

15 **PMIX_ATTR_UNDEF** `NULL` (`NULL`)
16 Constant representing an undefined attribute.

17 3.4.1 Initialization attributes

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

20 **PMIX_EVENT_BASE** `"pmix.evbase"` (`struct event_base *`)
21 Pointer to libevent¹ `event_base` to use in place of the internal progress thread.
22 **PMIX_SERVER_TOOL_SUPPORT** `"pmix.srvr.tool"` (`bool`)
23 The host RM wants to declare itself as willing to accept tool connection requests.
24 **PMIX_SERVER_REMOTE_CONNECTIONS** `"pmix.srvr.remote"` (`bool`)
25 Allow connections from remote tools. Forces the PMIx server to not exclusively use
26 loopback device.
27 **PMIX_SERVER_SYSTEM_SUPPORT** `"pmix.srvr.sys"` (`bool`)
28 The host RM wants to declare itself as being the local system server for PMIx connection
29 requests.
30 **PMIX_SERVER_TMPDIR** `"pmix.srvr.tmpdir"` (`char*`)
31 Top-level temporary directory for all *client* processes connected to this server, and where the
32 PMIx server will place its *tool* rendezvous point and contact information.
33 **PMIX_SYSTEM_TMPDIR** `"pmix.sys.tmpdir"` (`char*`)
34 Temporary directory for this system, and where a PMIx server that declares itself to be a
35 system-level server will place a *tool* rendezvous point and contact information.

¹<http://libevent.org/>

1 **PMIX_REGISTER_NODATA** "pmix.reg.nodata" (bool)
 2 Registration is for the namespace only. Do not copy job data.
 3 **PMIX_SERVER_ENABLE_MONITORING** "pmix.srv.monitor" (bool)
 4 Enable PMIx internal monitoring by the PMIx server.
 5 **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (char*)
 6 Name of the namespace to use for this PMIx server.
 7 **PMIX_SERVER_RANK** "pmix.srv.rank" (pmix_rank_t)
 8 Rank of this PMIx server

9 3.4.2 Tool-related attributes

10 These attributes are defined to assist PMIx-enabled tools to connect with the PMIx server by
 11 passing them into the **PMIx_tool_init** API - thus, they are not typically accessed via the
 12 **PMIx_Get** API.

13 **PMIX_TOOL_NAMESPACE** "pmix.tool.namespace" (char*)
 14 Name of the namespace to use for this tool.
 15 **PMIX_TOOL_RANK** "pmix.tool.rank" (uint32_t)
 16 Rank of this tool.
 17 **PMIX_SERVER_PIDINFO** "pmix.srvr.pidinfo" (pid_t)
 18 PID of the target PMIx server for a tool.
 19 **PMIX_CONNECT_TO_SYSTEM** "pmix.cnct.sys" (bool)
 20 The requestor requires that a connection be made only to a local, system-level PMIx server.
 21 **PMIX_CONNECT_SYSTEM_FIRST** "pmix.cnct.sys.first" (bool)
 22 Preferentially, look for a system-level PMIx server first.
 23 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)
 24 uniform resource identifier (URI) of the PMIx server to be contacted.
 25 **PMIX_SERVER_HOSTNAME** "pmix.srvr.host" (char*)
 26 Host where target PMIx server is located.
 27 **PMIX_CONNECT_MAX_RETRIES** "pmix.tool.mretries" (uint32_t)
 28 Maximum number of times to try to connect to PMIx server.
 29 **PMIX_CONNECT_RETRY_DELAY** "pmix.tool.retry" (uint32_t)
 30 Time in seconds between connection attempts to a PMIx server.
 31 **PMIX_TOOL_DO_NOT_CONNECT** "pmix.tool.nocon" (bool)
 32 The tool wants to use internal PMIx support, but does not want to connect to a PMIx server.

33 3.4.3 Identification attributes

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

36 **PMIX_USERID** "pmix.euid" (uint32_t)
 37 Effective user id.

1 **PMIX_GRPID** "pmix.egid" (uint32_t)
 2 Effective group id.
 3 **PMIX_DSTPATH** "pmix.dstpath" (char*)
 4 Path to shared memory data storage (dstore) files.
 5 **PMIX_VERSION_INFO** "pmix.version" (char*)
 6 PMIx version of contractor.
 7 **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)
 8 Programming model being initialized (e.g., "MPI" or "OpenMP")
 9 **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)
 10 Programming model implementation ID (e.g., "OpenMPI" or "MPICH")
 11 **PMIX_MODEL_LIBRARY_VERSION** "pmix.mld.vrs" (char*)
 12 Programming model version string (e.g., "2.1.1")
 13 **PMIX_THREADING_MODEL** "pmix.threads" (char*)
 14 Threading model used (e.g., "pthreads")
 15 **PMIX_REQUESTOR_IS_TOOL** "pmix.req.tool" (bool)
 16 The requesting process is a PMIx tool.
 17 **PMIX_REQUESTOR_IS_CLIENT** "pmix.req.client" (bool)
 18 The requesting process is a PMIx client.

19 3.4.4 UNIX socket rendezvous socket attributes

20 These attributes are used to describe a UNIX socket for rendezvous with the local RM by passing
 21 them into the relevant initialization API - thus, they are not typically accessed via the **PMIx_Get**
 22 API.

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

30 3.4.5 TCP connection attributes

31 These attributes are used to describe a TCP socket for rendezvous with the local RM by passing
 32 them into the relevant initialization API - thus, they are not typically accessed via the **PMIx_Get**
 33 API.

34 **PMIX_TCP_REPORT_URI** "pmix.tcp.repuri" (char*)
 35 If provided, directs that the TCP URI be reported and indicates the desired method of
 36 reporting: '-' for stdout, '+' for stderr, or filename.
 37 **PMIX_TCP_URI** "pmix.tcp.uri" (char*)

1 The URI of the PMIx server to connect to, or a file name containing it in the form of
2 `file:<name of file containing it>`.

3 **PMIX_TCP_IF_INCLUDE** `"pmix.tcp.ifinclude"` (`char*`)

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

6 **PMIX_TCP_IF_EXCLUDE** `"pmix.tcp.ifexclude"` (`char*`)

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

9 **PMIX_TCP_IPV4_PORT** `"pmix.tcp.ipv4"` (`int`)

10 The IPv4 port to be used.

11 **PMIX_TCP_IPV6_PORT** `"pmix.tcp.ipv6"` (`int`)

12 The IPv6 port to be used.

13 **PMIX_TCP_DISABLE_IPV4** `"pmix.tcp.disipv4"` (`bool`)

14 Set to `true` to disable IPv4 family of addresses.

15 **PMIX_TCP_DISABLE_IPV6** `"pmix.tcp.disipv6"` (`bool`)

16 Set to `true` to disable IPv6 family of addresses.

17 3.4.6 Global Data Storage (GDS) attributes

18 These attributes are used to define the behavior of the GDS used to manage key/value pairs by
19 passing them into the relevant initialization API - thus, they are not typically accessed via the
20 **PMIx_Get** API.

21 **PMIX_GDS_MODULE** `"pmix.gds.mod"` (`char*`)

22 Comma-delimited string of desired modules.

23 3.4.7 General process-level attributes

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

25 **PMIX_CPUSSET** `"pmix.cpuset"` (`char*`)

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

27 **PMIX_CREDENTIAL** `"pmix.cred"` (`char*`)

28 Security credential assigned to the process.

29 **PMIX_SPAWNED** `"pmix.spawned"` (`bool`)

30 `true` if this process resulted from a call to **PMIx_Spawn**.

31 **PMIX_ARCH** `"pmix.arch"` (`uint32_t`)

32 Architecture flag.

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

1 3.4.8 Scratch directory attributes

2 These attributes are used to define an application scratch directory and are referenced using the
3 [PMIX_RANK_WILDCARD](#) rank.

4 **PMIX_TMPDIR** "pmix.tmpdir" (char*)

5 Full path to the top-level temporary directory assigned to the session.

6 **PMIX_NSDIR** "pmix.nmdir" (char*)

7 Full path to the temporary directory assigned to the namespace, under [PMIX_TMPDIR](#) .

8 **PMIX_PROCDIR** "pmix.pdir" (char*)

9 Full path to the subdirectory under [PMIX_NSDIR](#) assigned to the process.

10 **PMIX_TDIR_RMCLEAN** "pmix.tdir.rmclean" (bool)

11 Resource Manager will clean session directories

12 3.4.9 Relative Rank Descriptive Attributes

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

15 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)

16 Process identifier

17 **PMIX_NAMESPACE** "pmix.namespace" (char*)

18 Namespace of the job.

19 **PMIX_JOBID** "pmix.jobid" (char*)

20 Job identifier assigned by the scheduler.

21 **PMIX_APPNUM** "pmix.appnum" (uint32_t)

22 Application number within the job.

23 **PMIX_RANK** "pmix.rank" (pmix_rank_t)

24 Process rank within the job.

25 **PMIX_GLOBAL_RANK** "pmix.grank" (pmix_rank_t)

26 Process rank spanning across all jobs in this session.

27 **PMIX_APP_RANK** "pmix.apprank" (pmix_rank_t)

28 Process rank within this application.

29 **PMIX_NPROC_OFFSET** "pmix.offset" (pmix_rank_t)

30 Starting global rank of this job - referenced using [PMIX_RANK_WILDCARD](#) .

31 **PMIX_LOCAL_RANK** "pmix.lrank" (uint16_t)

32 Local rank on this node within this job.

33 **PMIX_NODE_RANK** "pmix.nrank" (uint16_t)

34 Process rank on this node spanning all jobs.

35 **PMIX_LOCALLDR** "pmix.lldr" (pmix_rank_t)

36 Lowest rank on this node within this job - referenced using [PMIX_RANK_WILDCARD](#) .

37 **PMIX_APPLDR** "pmix.aldr" (pmix_rank_t)

38 Lowest rank in this application within this job - referenced using [PMIX_RANK_WILDCARD](#) .

39 **PMIX_PROC_PID** "pmix.ppid" (pid_t)

1 PID of specified process.

2 **PMIX_SESSION_ID** "pmix.session.id" (uint32_t)

3 Session identifier - referenced using [PMIX_RANK_WILDCARD](#) .

4 **PMIX_NODE_LIST** "pmix.nlist" (char*)

5 Comma-delimited list of nodes running processes for the specified namespace - referenced

6 using [PMIX_RANK_WILDCARD](#) .

7 **PMIX_ALLOCATED_NODELIST** "pmix.alist" (char*)

8 Comma-delimited list of all nodes in this allocation regardless of whether or not they

9 currently host processes - referenced using [PMIX_RANK_WILDCARD](#) .

10 **PMIX_HOSTNAME** "pmix.hname" (char*)

11 Name of the host where the specified process is running.

12 **PMIX_NODEID** "pmix.nodeid" (uint32_t)

13 Node identifier where the specified process is located, expressed as the node's index

14 (beginning at zero) in the array resulting from expansion of the [PMIX_NODE_MAP](#) regular

15 expression for the [job](#)

16 **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)

17 Comma-delimited list of ranks on this node within the specified namespace - referenced

18 using [PMIX_RANK_WILDCARD](#) .

19 **PMIX_LOCAL_PROCS** "pmix.lprocs" (pmix_proc_t array)

20 Array of [pmix_proc_t](#) of all processes on the specified node - referenced using

21 [PMIX_RANK_WILDCARD](#) .

22 **PMIX_LOCAL_CPUSSETS** "pmix.lcpus" (char*)

23 Colon-delimited cpusets of local peers within the specified namespace - referenced using

24 [PMIX_RANK_WILDCARD](#) .

25 **PMIX_PROC_URI** "pmix.puri" (char*)

26 URI containing contact information for a given process.

27 **PMIX_LOCALITY** "pmix.loc" (uint16_t)

28 Relative locality of the specified process to the requestor.

29 **PMIX_PARENT_ID** "pmix.parent" (pmix_proc_t)

30 Process identifier of the parent process of the calling process.

31 **3.4.10 Information retrieval attributes**

32 The following attributes are used to specify the level of information (e.g., [session](#) , [job](#) , or

33 [application](#)) being requested where ambiguity may exist - see [5.1.5](#) for examples of their use.

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

35 Return information about the specified session. If information about a session other than the

36 one containing the requesting process is desired, then the attribute array must contain a

37 [PMIX_SESSION_ID](#) attribute identifying the desired target.

38 **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
11 is requested about an application in a job or session other than the one containing the
12 requesting 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.

17 3.4.11 Information storage attributes

18 The following attributes are used to assemble information by its level (e.g., **session** , **job** , or
19 **application**) for storage where ambiguity may exist - see 10.1.3.1 for examples of their use.

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

21 Provide an array of **pmix_info_t** containing session-level information. The
22 **PMIX_SESSION_ID** attribute is *required* to be included in the array.

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

24 Provide an array of **pmix_info_t** containing job-level information. Information is
25 registered one job (aka namespace) at a time via the **PMIx_server_register_namespace**
26 API. Thus, there is no requirement that the array contain either the **PMIX_NAMESPACE** or
27 **PMIX_JOBID** attributes, though either or both of them *may* be included.

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

29 Provide an array of **pmix_info_t** containing app-level information. The **PMIX_NAMESPACE**
30 or **PMIX_JOBID** attributes of the **job** containing the application, plus its
31 **PMIX_APPNUM** attribute, are *required* to be included in the array.

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

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

36 3.4.12 Size information attributes

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

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

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

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

6 Total number of processes in this job across all contained applications

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

8 Number of applications in this job.

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

10 Number of processes in this application.

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

12 Number of processes in this job on this node.

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

14 Number of processes across all jobs on this node.

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

16 Maximum number of processes for this job.

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

18 Number of nodes in this session or namespace.

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

20 Number of slots allocated to the session, namespace, or application.

21 3.4.13 Memory information attributes

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

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

25 Total available physical memory on this node.

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

27 Megabytes of memory currently used by the RM daemon.

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

29 Average Megabytes of memory used by client processes.

30 3.4.14 Topology information attributes

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

33 **PMIX_NET_TOPO** "pmix.ntopo" (char*)

34 eXtensible Markup Language (XML) representation of the network topology.

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

36 XML representation of local node topology.

37 **PMIX_NODE_LIST** "pmix.nlist" (char*)

1 Comma-delimited list of nodes running processes for this job.

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

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

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

5 Topology signature string.

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

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

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

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

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

16 This means numa node 0, socket 0, L3 caches 0, 1, 2, 3, 4, L2 caches 0-4, L1 caches
17 0-4, cores 0, 1, 2, 3, 4, and hardware threads 0-39.

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

19 Address of the hwloc shared memory segment.

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

21 Size of the hwloc shared memory segment.

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

23 Path to the hwloc shared memory file.

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

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

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

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

28 3.4.15 Request-related attributes

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

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

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

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

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

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

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

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

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

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

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

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

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

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

11 Notify the parent process upon termination of child job.

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

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

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

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

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

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

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_EMBED_BARRIER** "pmix.embed.barrier" (bool)

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

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

26 Status to be returned upon job termination.

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

28 Process state

29 3.4.16 Server-to-PMIx library attributes

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

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

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

36 **PMIX_PROC_DATA** "pmix.pdata" (pmix_data_array_t)

37 Array of process data. Starts with rank, then contains more data.

38 **PMIX_NODE_MAP** "pmix.nmap" (char*)

39 Regular expression of nodes - see [10.1.3.1](#) for an explanation of its generation.

40 **PMIX_PROC_MAP** "pmix.pmap" (char*)

1 Regular expression describing processes on each node - see [10.1.3.1](#) for an explanation of its
 2 generation.
 3 **PMIX_ANL_MAP** "pmix.anlmap" (char*)
 4 Process mapping in Argonne National Laboratory's PMI-1/PMI-2 notation.
 5 **PMIX_APP_MAP_TYPE** "pmix.apmap.type" (char*)
 6 Type of mapping used to layout the application (e.g., **cyclic**).
 7 **PMIX_APP_MAP_REGEX** "pmix.apmap.regex" (char*)
 8 Regular expression describing the result of the process mapping.

9 3.4.17 Server-to-Client attributes

10 Attributes used internally to communicate data from the PMIx server to the PMIx client - they do
 11 not represent values accessed using the **PMIx_Get** API.
 12 **PMIX_PROC_BLOB** "pmix.pblob" (pmix_byte_object_t)
 13 Packed blob of process data.
 14 **PMIX_MAP_BLOB** "pmix.mblob" (pmix_byte_object_t)
 15 Packed blob of process location.

16 3.4.18 Event handler registration and notification attributes

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

▼ Advice to users ▼

19 The event handler subsystem defined in the PMIx *ad hoc* version 1 Standard was completely
 20 overhauled in version 2 to resolve design flaws. Deprecated attributes shown below were therefore
 21 removed in the version 2 Standard.

22 **PMIX_ERROR_NAME** "pmix.errname" (pmix_status_t)
 23 Specific error to be notified
 24 **PMIX_ERROR_GROUP_COMM** "pmix.errgroup.comm" (bool)
 25 Set true to get comm errors notification
 26 **PMIX_ERROR_GROUP_ABORT** "pmix.errgroup.abort" (bool)
 27 Set true to get abort errors notification
 28 **PMIX_ERROR_GROUP_MIGRATE** "pmix.errgroup.migrate" (bool)
 29 Set true to get migrate errors notification
 30 **PMIX_ERROR_GROUP_RESOURCE** "pmix.errgroup.resource" (bool)
 31 Set true to get resource errors notification
 32 **PMIX_ERROR_GROUP_SPAWN** "pmix.errgroup.spawn" (bool)
 33 Set true to get spawn errors notification
 34 **PMIX_ERROR_GROUP_NODE** "pmix.errgroup.node" (bool)

1 Set true to get node status notification

2 **PMIX_ERROR_GROUP_LOCAL** "pmix.errgroup.local" (bool)

3 Set true to get local errors notification

4 **PMIX_ERROR_GROUP_GENERAL** "pmix.errgroup.gen" (bool)

5 Set true to get notified of generic errors

6 **PMIX_ERROR_HANDLER_ID** "pmix.errhandler.id" (int)

7 Errhandler reference id of notification being reported

8 **PMIX_EVENT_HDLR_NAME** "pmix.evname" (char*)

9 String name identifying this handler.

10 **PMIX_EVENT_HDLR_FIRST** "pmix.evfirst" (bool)

11 Invoke this event handler before any other handlers.

12 **PMIX_EVENT_HDLR_LAST** "pmix.evlast" (bool)

13 Invoke this event handler after all other handlers have been called.

14 **PMIX_EVENT_HDLR_FIRST_IN_CATEGORY** "pmix.evfirstcat" (bool)

15 Invoke this event handler before any other handlers in this category.

16 **PMIX_EVENT_HDLR_LAST_IN_CATEGORY** "pmix.evlastcat" (bool)

17 Invoke this event handler after all other handlers in this category have been called.

18 **PMIX_EVENT_HDLR_BEFORE** "pmix.evbefore" (char*)

19 Put this event handler immediately before the one specified in the (char*) value.

20 **PMIX_EVENT_HDLR_AFTER** "pmix.evafter" (char*)

21 Put this event handler immediately after the one specified in the (char*) value.

22 **PMIX_EVENT_HDLR_PREPEND** "pmix.evprepend" (bool)

23 Prepend this handler to the precedence list within its category.

24 **PMIX_EVENT_HDLR_APPEND** "pmix.evappend" (bool)

25 Append this handler to the precedence list within its category.

26 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrange" (pmix_data_array_t*)

27 Array of `pmix_proc_t` defining range of event notification.

28 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)

29 The single process that was affected.

30 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (pmix_data_array_t*)

31 Array of `pmix_proc_t` defining affected processes.

32 **PMIX_EVENT_NON_DEFAULT** "pmix.evnondef" (bool)

33 Event is not to be delivered to default event handlers.

34 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (void *)

35 Object to be returned whenever the registered callback function `cbfunc` is invoked. The

36 object will *only* be returned to the process that registered it.

37 **PMIX_EVENT_DO_NOT_CACHE** "pmix.evnocache" (bool)

38 Instruct the PMIx server not to cache the event.

39 **PMIX_EVENT_SILENT_TERMINATION** "pmix.avsilentterm" (bool)

40 Do not generate an event when this job normally terminates.

1 3.4.19 Fault tolerance attributes

2 Attributes to support fault tolerance behaviors - they are values passed to the event notification API
3 and therefore are not accessed using the [PMIx_Get](#) API.

4 **PMIX_EVENT_TERMINATE_SESSION** "pmix.evterm.sess" (bool)

5 The RM intends to terminate this session.

6 **PMIX_EVENT_TERMINATE_JOB** "pmix.evterm.job" (bool)

7 The RM intends to terminate this job.

8 **PMIX_EVENT_TERMINATE_NODE** "pmix.evterm.node" (bool)

9 The RM intends to terminate all processes on this node.

10 **PMIX_EVENT_TERMINATE_PROC** "pmix.evterm.proc" (bool)

11 The RM intends to terminate just this process.

12 **PMIX_EVENT_ACTION_TIMEOUT** "pmix.evtimeout" (int)

13 The time in seconds before the RM will execute error response.

14 **PMIX_EVENT_NO_TERMINATION** "pmix.evnoterm" (bool)

15 Indicates that the handler has satisfactorily handled the event and believes termination of the
16 application is not required.

17 **PMIX_EVENT_WANT_TERMINATION** "pmix.evterm" (bool)

18 Indicates that the handler has determined that the application should be terminated

19 3.4.20 Spawn attributes

20 Attributes used to describe [PMIx_Spawn](#) behavior - they are values passed to the [PMIx_Spawn](#)
21 API and therefore are not accessed using the [PMIx_Get](#) API when used in that context. However,
22 some of the attributes defined in this section can be provided by the host environment for other
23 purposes - e.g., the environment might provide the [PMIX_MAPPER](#) attribute in the job-related
24 information so that an application can use [PMIx_Get](#) to discover the layout algorithm used for
25 determining process locations. Multi-use attributes and their respective access reference rank are
26 denoted below.

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

28 Name of personality to use.

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

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

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

32 Hostfile to use for spawned processes.

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

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

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

36 Hostfile listing hosts to add to existing allocation.

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

38 Prefix to use for starting spawned processes.

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

1 Working directory for spawned processes.

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

3 Mapping mechanism to use for placing spawned processes - when accessed using

4 **PMIx_Get** , use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping

5 mechanism used for the provided namespace.

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

7 Display process mapping upon spawn.

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

9 Number of processes to spawn on each identified resource.

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

11 Process mapping policy - when accessed using **PMIx_Get** , use the

12 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the

13 provided namespace

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

15 Process ranking policy - when accessed using **PMIx_Get** , use the

16 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the

17 provided namespace

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

19 Process binding policy - when accessed using **PMIx_Get** , use the

20 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the

21 provided namespace

22 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

23 Preload binaries onto nodes.

24 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

25 Comma-delimited list of files to pre-position on nodes.

26 **PMIX_NON_PMI** "pmix.nonpmi" (bool)

27 Spawned processes will not call **PMIx_Init** .

28 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)

29 Spawned process rank that is to receive **stdin**.

30 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (bool)

31 Forward this process's **stdin** to the designated process.

32 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)

33 Forward **stdout** from spawned processes to this process.

34 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)

35 Forward **stderr** from spawned processes to this process.

36 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)

37 Spawned application consists of debugger daemons.

38 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)

39 Designated application is to be spawned as a disconnected job. Meaning that it is not part of

40 the "comm_world" of the parent process.

41 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

1 Set the application's current working directory to the session working directory assigned by
2 the RM - when accessed using `PMIx_Get`, use the `PMIX_RANK_WILDCARD` value for
3 the rank to discover the session working directory assigned to the provided namespace

4 `PMIX_TAG_OUTPUT` "pmix.tagout" (bool)

5 Tag application output with the identity of the source process.

6 `PMIX_TIMESTAMP_OUTPUT` "pmix.tsout" (bool)

7 Timestamp output from applications.

8 `PMIX_MERGE_STDERR_STDOUT` "pmix.mergeerrout" (bool)

9 Merge `stdout` and `stderr` streams from application processes.

10 `PMIX_OUTPUT_TO_FILE` "pmix.outfile" (char*)

11 Output application output to the specified file.

12 `PMIX_INDEX_ARGV` "pmix.indxargv" (bool)

13 Mark the `argv` with the rank of the process.

14 `PMIX_CPUS_PER_PROC` "pmix.cpusperproc" (uint32_t)

15 Number of cpus to assign to each rank - when accessed using `PMIx_Get`, use the
16 `PMIX_RANK_WILDCARD` value for the rank to discover the cpus/process assigned to the
17 provided namespace

18 `PMIX_NO_PROCS_ON_HEAD` "pmix.nolocal" (bool)

19 Do not place processes on the head node.

20 `PMIX_NO_OVERSUBSCRIBE` "pmix.noover" (bool)

21 Do not oversubscribe the cpus.

22 `PMIX_REPORT_BINDINGS` "pmix.repbinding" (bool)

23 Report bindings of the individual processes.

24 `PMIX_CPU_LIST` "pmix.cpulist" (char*)

25 List of cpus to use for this job - when accessed using `PMIx_Get`, use the
26 `PMIX_RANK_WILDCARD` value for the rank to discover the cpu list used for the provided
27 namespace

28 `PMIX_JOB_RECOVERABLE` "pmix.recover" (bool)

29 Application supports recoverable operations.

30 `PMIX_JOB_CONTINUOUS` "pmix.continuous" (bool)

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

32 `PMIX_MAX_RESTARTS` "pmix.maxrestarts" (uint32_t)

33 Maximum number of times to restart a job - when accessed using `PMIx_Get`, use the
34 `PMIX_RANK_WILDCARD` value for the rank to discover the max restarts for the provided
35 namespace

36 3.4.21 Query attributes

37 Attributes used to describe `PMIx_Query_info_nb` behavior - these are values passed to the
38 `PMIx_Query_info_nb` API and therefore are not passed to the `PMIx_Get` API.

39 `PMIX_QUERY_REFRESH_CACHE` "pmix.qry.rfsh" (bool)

40 Retrieve updated information from server.

1 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (char*)
 2 Request a comma-delimited list of active namespaces.

3 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)
 4 Status of a specified, currently executing job.

5 **PMIX_QUERY_QUEUE_LIST** "pmix.qry.qlst" (char*)
 6 Request a comma-delimited list of scheduler queues.

7 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (TBD)
 8 Status of a specified scheduler queue.

9 **PMIX_QUERY_PROC_TABLE** "pmix.qry.ptable" (char*)
 10 Input namespace of the job whose information is being requested returns (
 11 **pmix_data_array_t**) an array of **pmix_proc_info_t**.

12 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lptable" (char*)
 13 Input namespace of the job whose information is being requested returns (
 14 **pmix_data_array_t**) an array of **pmix_proc_info_t** for processes in job on same
 15 node.

16 **PMIX_QUERY_LOCAL_ONLY** "pmix.qry.local" (bool)
 17 Constrain the query to local information only.

18 **PMIX_QUERY_AUTHORIZATIONS** "pmix.qry.auths" (bool)
 19 Return operations the PMIx tool is authorized to perform.

20 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)
 21 Return a comma-delimited list of supported spawn attributes.

22 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)
 23 Return a comma-delimited list of supported debug attributes.

24 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)
 25 Return information on memory usage for the processes indicated in the qualifiers.

26 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
 27 Report average values.

28 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
 29 Report minimum and maximum values.

30 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 31 String identifier of the allocation whose status is being requested.

32 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
 33 Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.

34 3.4.22 Log attributes

35 Attributes used to describe **PMIx_Log_nb** behavior - these are values passed to the
 36 **PMIx_Log_nb** API and therefore are not accessed using the **PMIx_Get** API.

37 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)
 38 Log string to **stderr**.

39 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)
 40 Log string to **stdout**.

1 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)
 2 Log data to syslog. Defaults to **ERROR** priority.
 3 **PMIX_LOG_MSG** "pmix.log.msg" (pmix_byte_object_t)
 4 Message blob to be sent somewhere.
 5 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)
 6 Log via email based on **pmix_info_t** containing directives.
 7 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)
 8 Comma-delimited list of email addresses that are to receive the message.
 9 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)
 10 Subject line for email.
 11 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)
 12 Message to be included in email.

13 3.4.23 Debugger attributes

14 Attributes used to assist debuggers - these are values that can be passed to the **PMIx_Spawn** or
 15 **PMIx_Init** APIs. Some may be accessed using the **PMIx_Get** API with the
 16 **PMIX_RANK_WILDCARD** rank.

17 ===== **PMIX_DEBUG_STOP_ON_EXEC** "pmix.dbg.exec" (bool)
 18 Passed to **PMIx_Spawn** to indicate that the specified application is being spawned under
 19 debugger, and that the launcher is to pause the resulting application processes on first
 20 instruction for debugger attach.
 21 **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)
 22 Passed to **PMIx_Spawn** to indicate that the specified application is being spawned under
 23 debugger, and that the PMIx client library is to pause the resulting application processes
 24 during **PMIx_Init** until debugger attach and release.
 25 **PMIX_DEBUG_WAIT_FOR_NOTIFY** "pmix.dbg.notify" (bool)
 26 Passed to **PMIx_Spawn** to indicate that the specified application is being spawned under
 27 debugger, and that the resulting application processes are to pause at some
 28 application-determined location until debugger attach and release.
 29 **PMIX_DEBUG_JOB** "pmix.dbg.job" (char*)
 30 Namespace of the job to be debugged - provided to the debugger upon launch.
 31 **PMIX_DEBUG_WAITING_FOR_NOTIFY** "pmix.dbg.waiting" (bool)
 32 Job to be debugged is waiting for a release - this is not a value accessed using the
 33 **PMIx_Get** API.

34 3.4.24 Resource manager attributes

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

39 **PMIX_RM_NAME** "pmix.rm.name" (char*)

1 String name of the RM.
2 **PMIX_RM_VERSION** "pmix.rm.version" (char*)
3 RM version string.

4 3.4.25 Environment variable attributes

5 Attributes used to adjust environment variables - these are values passed to the **PMIx_Spawn** API
6 and are not accessed using the **PMIx_Get** API.

7 **PMIX_SET_ENVAR** "pmix.set.envar" (char*)
8 String "key=value" value shall be put into the environment.
9 **PMIX_UNSET_ENVAR** "pmix.unset.envar" (char*)
10 Unset the environment variable specified in the string.

11 3.4.26 Job Allocation attributes

12 Attributes used to describe the job allocation - these are values passed to the
13 **PMIx_Allocation_request_nb** API and are not accessed using the **PMIx_Get** API

14 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)
15 Provide a string identifier for this allocation request which can later be used to query status
16 of the request.
17 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)
18 The number of nodes.
19 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
20 Regular expression of the specific nodes.
21 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
22 Number of cpus.
23 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
24 Regular expression of the number of cpus for each node.
25 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
26 Regular expression of the specific cpus indicating the cpus involved.
27 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
28 Number of Megabytes.
29 **PMIX_ALLOC_NETWORK** "pmix.alloc.net" (array)
30 Array of **pmix_info_t** describing requested network resources. If not given as part of an
31 **pmix_info_t** struct that identifies the involved nodes, then the description will be
32 applied across all nodes in the requestor's allocation.
33 **PMIX_ALLOC_NETWORK_ID** "pmix.alloc.netid" (char*)
34 Name of the network.
35 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)
36 Mbits/sec.
37 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)

1 Quality of service level.
2 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
3 Time in seconds.

4 **3.4.27 Job control attributes**

5 Attributes used to request control operations on an executing application - these are values passed
6 to the **PMIx_Job_control_nb** API and are not accessed using the **PMIx_Get** API.

7 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)
8 Provide a string identifier for this request.
9 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)
10 Pause the specified processes.
11 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)
12 Resume ("un-pause") the specified processes.
13 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
14 Cancel the specified request (**NULL** implies cancel all requests from this requestor).
15 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
16 Forcibly terminate the specified processes and cleanup.
17 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
18 Restart the specified processes using the given checkpoint ID.
19 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
20 Checkpoint the specified processes and assign the given ID to it.
21 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
22 Use event notification to trigger a process checkpoint.
23 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
24 Use the given signal to trigger a process checkpoint.
25 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
26 Time in seconds to wait for a checkpoint to complete.
27 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
28 "pmix.jctrl.ckmethod" (pmix_data_array_t)
29 Array of **pmix_info_t** declaring each method and value supported by this application.
30 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
31 Send given signal to specified processes.
32 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
33 Regular expression identifying nodes that are to be provisioned.
34 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
35 Name of the image that is to be provisioned.
36 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
37 Indicate that the job can be pre-empted.
38 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
39 Politely terminate the specified processes.

1 3.4.28 Monitoring attributes

2 Attributes used to control monitoring of an executing application- these are values passed to the
3 `PMIx_Process_monitor_nb` API and are not accessed using the `PMIx_Get` API.

4 `PMIX_MONITOR_ID` `"pmix.monitor.id"` (`char*`)
5 Provide a string identifier for this request.
6 `PMIX_MONITOR_CANCEL` `"pmix.monitor.cancel"` (`char*`)
7 Identifier to be canceled (`NULL` means cancel all monitoring for this process).
8 `PMIX_MONITOR_APP_CONTROL` `"pmix.monitor.appctrl"` (`bool`)
9 The application desires to control the response to a monitoring event.
10 `PMIX_MONITOR_HEARTBEAT` `"pmix.monitor.mbeat"` (`void`)
11 Register to have the PMIx server monitor the requestor for heartbeats.
12 `PMIX_SEND_HEARTBEAT` `"pmix.monitor.beat"` (`void`)
13 Send heartbeat to local PMIx server.
14 `PMIX_MONITOR_HEARTBEAT_TIME` `"pmix.monitor.btime"` (`uint32_t`)
15 Time in seconds before declaring heartbeat missed.
16 `PMIX_MONITOR_HEARTBEAT_DROPS` `"pmix.monitor.bdrop"` (`uint32_t`)
17 Number of heartbeats that can be missed before generating the event.
18 `PMIX_MONITOR_FILE` `"pmix.monitor.fmon"` (`char*`)
19 Register to monitor file for signs of life.
20 `PMIX_MONITOR_FILE_SIZE` `"pmix.monitor.fsize"` (`bool`)
21 Monitor size of given file is growing to determine if the application is running.
22 `PMIX_MONITOR_FILE_ACCESS` `"pmix.monitor.faccess"` (`char*`)
23 Monitor time since last access of given file to determine if the application is running.
24 `PMIX_MONITOR_FILE_MODIFY` `"pmix.monitor.fmod"` (`char*`)
25 Monitor time since last modified of given file to determine if the application is running.
26 `PMIX_MONITOR_FILE_CHECK_TIME` `"pmix.monitor.ftime"` (`uint32_t`)
27 Time in seconds between checking the file.
28 `PMIX_MONITOR_FILE_DROPS` `"pmix.monitor.fdrop"` (`uint32_t`)
29 Number of file checks that can be missed before generating the event.

30 3.5 Callback Functions

31 PMIx provides blocking and nonblocking versions of most APIs. In the nonblocking versions, a
32 callback is activated upon completion of the the operation. This section describes many of those
33 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_fence_nb_fn_t` and
17 `pmix_server_dmodex_req_fn_t` PMIx server operations to return modex BCX data.

PMIx v1.0

```
▼ C ▼  
18 typedef void (*pmix_modex_cbfunc_t)  
19     (pmix_status_t status,  
20      const char *data, size_t ndata,  
21      void *cbdata,  
22      pmix_release_cbfunc_t release_fn,  
23      void *release_cbdata)  
▲ C ▲
```

24 IN `status`

25 Status associated with the operation (handle)

26 IN `data`

27 Data to be passed (pointer)

1 **IN** **ndata**
 2 size of the data (**size_t**)
 3 **IN** **cbdata**
 4 Callback data passed to original API call (memory reference)
 5 **IN** **release_fn**
 6 Callback for releasing *data* (function pointer)
 7 **IN** **release_cbdata**
 8 Pointer to be passed to *release_fn* (memory reference)

9 **Description**

10 A callback function that is solely used by PMIx servers, and not clients, to return modex BCX data
 11 in response to “fence” and “get” operations. The returned blob contains the data collected from
 12 each server participating in the operation.

13 **3.5.3 Spawn Callback Function**

14 **Summary**

15 The `pmix_spawn_cbfunc_t` is used on the PMIx client side by `PMIx_Spawn_nb` and on
 16 the PMIx server side by `pmix_server_spawn_fn_t`.

17 *PMIx v1.0*

```

17       typedef void (*pmix_spawn_cbfunc_t)
18           (pmix_status_t status,
19            pmix_namespace_t nspace, void *cbdata);

```

20 **IN** **status**
 21 Status associated with the operation (handle)
 22 **IN** **namespace**
 23 Namespace string (`pmix_namespace_t`)
 24 **IN** **cbdata**
 25 Callback data passed to original API call (memory reference)

26 **Description**

27 The callback will be executed upon launch of the specified applications in `PMIx_Spawn_nb`, or
 28 upon failure to launch any of them.

29 The *status* of the callback will indicate whether or not the spawn succeeded. The *namespace* of the
 30 spawned processes will be returned, along with any provided callback data. Note that the returned
 31 *namespace* value will not be protected by the PRI upon return from the callback function, so the
 32 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

```
4 typedef void (*pmix_op_cbfunc_t)
5     (pmix_status_t status, void *cbdata);
```

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

```
17 typedef void (*pmix_lookup_cbfunc_t)
18     (pmix_status_t status,
19      pmix_pdata_t data[], size_t ndata,
20      void *cbdata);
```

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

```

1 typedef void (*pmix_info_cbfunc_t)
2     (pmix_status_t status,
3      pmix_info_t info[], size_t ninfo,
4      void *cbdata,
5      pmix_release_cbfunc_t release_fn,
6      void *release_cbdata);

```

IN status
Status associated with the operation (`pmix_status_t`)

IN info
Array of `pmix_info_t` returned by the operation (pointer)

IN ninfo
Number of elements in the *info* array (`size_t`)

IN cbdata
Callback data passed to original API call (memory reference)

IN release_fn
Function to be called when done with the *info* data (function pointer)

IN release_cbdata
Callback data to be passed to *release_fn* (memory reference)

Description

The *status* indicates if requested data was found or not. An array of `pmix_info_t` will contain the key/value pairs.

3.5.8 Event Handler Registration Callback Function

The `pmix_evhdlr_reg_cbfunc_t` callback function.

Advice to users

The PMIx *ad hoc* v1.0 Standard defined an error handler registration callback function with a compatible signature, but with a different type definition function name (`pmix_errhandler_reg_cbfunc_t`). It was removed from the v2.0 Standard and is not included in this document to avoid confusion.

PMIx v2.0

```

1 typedef void (*pmix_evhdlr_reg_cbfunc_t)
2     (pmix_status_t status,
3      size_t evhdlr_ref,
4      void *cbdata)

```

```

5 IN  status
6     Status indicates if the request was successful or not (pmix\_status\_t)
7 IN  evhdlr_ref
8     Reference assigned to the event handler by PMIx — this reference * must be used to
9     deregister the err handler (size\_t)
10 IN  cbdata
11     Callback data passed to original API call (memory reference)

```

Description

Define a callback function for calls to [PMIx_Register_event_handler](#)

3.5.9 Notification Handler Completion Callback Function

Summary

The [pmix_event_notification_cbfunc_fn_t](#) is called by event handlers to indicate completion of their operations.

PMIx v2.0

```

18 typedef void (*pmix_event_notification_cbfunc_fn_t)
19     (pmix_status_t status,
20      pmix_info_t *results, size_t nresults,
21      pmix_op_cbfunc_t cbfunc, void *thiscbdata,
22      void *notification_cbdata);

```

```

23 IN  status
24     Status returned by the event handler's operation (pmix\_status\_t)
25 IN  results
26     Results from this event handler's operation on the event (pmix\_info\_t)
27 IN  nresults
28     Number of elements in the results array (size\_t)
29 IN  cbfunc
30     pmix\_op\_cbfunc\_t function to be executed when PMIx completes processing the
31     callback (function reference)

```

1 **IN** **thiscbdata**
2 Callback data that was passed in to the handler (memory reference)
3 **IN** **cbdata**
4 Callback data to be returned when PMIx executes cbfunc (memory reference)

5 **Description**

6 Define a callback by which an event handler can notify the PMIx library that it has completed its
7 response to the notification. The handler is *required* to execute this callback so the library can
8 determine if additional handlers need to be called. The handler shall return
9 [PMIX_ERR_EVENT_COMPLETE](#) if no further action is required. The return status of each event
10 handler and any returned [pmix_info_t](#) structures will be added to the *results* array of
11 [pmix_info_t](#) passed to any subsequent event handlers to help guide their operation.
12 If non-NULL, the provided callback function will be called to allow the event handler to release the
13 provided info array and execute any other required cleanup operations.

14 **3.5.10 Notification Function**

15 **Summary**

16 The [pmix_notification_fn_t](#) is called by PMIx to deliver notification of an event.

▼ Advice to users ▼

17 The PMIx *ad hoc* v1.0 Standard defined an error notification function with an identical name, but
18 different signature than the v2.0 Standard described below. The *ad hoc* v1.0 version was removed
19 from the v2.0 Standard is not included in this document to avoid confusion.



PMIx v2.0

C

```
20      typedef void (*pmix_notification_fn_t)  
21           (size_t evhdlr_registration_id,  
22            pmix_status_t status,  
23            const pmix_proc_t *source,  
24            pmix_info_t info[], size_t ninfo,  
25            pmix_info_t results[], size_t nresults,  
26            pmix_event_notification_cbfunc_fn_t cbfunc,  
27            void *cbdata);
```

1 **IN evhdlr_registration_id**
 2 Registration number of the handler being called (**size_t**)
 3 **IN status**
 4 Status associated with the operation (**pmix_status_t**)
 5 **IN source**
 6 Identifier of the process that generated the event (**pmix_proc_t**). If the source is the
 7 SMS, then the nspace will be empty and the rank will be PMIX_RANK_UNDEF
 8 **IN info**
 9 Information describing the event (**pmix_info_t**). This argument will be **NULL** if no
 10 additional information was provided by the event generator.
 11 **IN ninfo**
 12 Number of elements in the info array (**size_t**)
 13 **IN results**
 14 Aggregated results from prior event handlers servicing this event (**pmix_info_t**). This
 15 argument will be **NULL** if this is the first handler servicing the event, or if no prior handlers
 16 provided results.
 17 **IN nresults**
 18 Number of elements in the results array (**size_t**)
 19 **IN cbfunc**
 20 **pmix_event_notification_cbfunc_fn_t** callback function to be executed upon
 21 completion of the handler's operation and prior to handler return (function reference).
 22 **IN cbdata**
 23 Callback data to be passed to cbfunc (memory reference)

24 Description

25 Note that different RMs may provide differing levels of support for event notification to application
 26 processes. Thus, the *info* array may be **NULL** or may contain detailed information of the event. It is
 27 the responsibility of the application to parse any provided info array for defined key-values if it so
 28 desires.

▼ Advice to users ▼

29 Possible uses of the *info* array include:

- 30 • for the host RM to alert the process as to planned actions, such as aborting the session, in
 31 response to the reported event
- 32 • provide a timeout for alternative action to occur, such as for the application to request an
 33 alternate response to the event

1 For example, the RM might alert the application to the failure of a node that resulted in termination
2 of several processes, and indicate that the overall session will be aborted unless the application
3 requests an alternative behavior in the next 5 seconds. The application then has time to respond
4 with a checkpoint request, or a request to recover from the failure by obtaining replacement nodes
5 and restarting from some earlier checkpoint.

6 Support for these options is left to the discretion of the host RM. Info keys are included in the
7 common definitions above but may be augmented by environment vendors.



Advice to PMIx server hosts

8 On the server side, the notification function is used to inform the PMIx server library's host of a
9 detected event in the PMIx server library. Events generated by PMIx clients are communicated to
10 the PMIx server library, but will be relayed to the host via the
11 [pmix_server_notify_event_fn_t](#) function pointer, if provided.



12 3.5.11 Server Setup Application Callback Function

13 The [PMIx_server_setup_application](#) callback function.

14 Summary

15 Provide a function by which the resource manager can receive application-specific environmental
16 variables and other setup data prior to launch of an application.

1 **Format**

PMIx v2.0

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
18 reference)
- 19 **IN cbdata**
20 Data to be passed to the *cbfunc* callback function (memory reference)

21 **Description**

22 Define a function to be called by the PMIx server library for return of application-specific setup
23 data in response to a request from the host RM. The returned *info* array is owned by the PMIx
24 server library and will be free'd when the provided *cbfunc* is called.

25 **3.5.12 Server Direct Modex Response Callback Function**

26 The [PMIx_server_dmodex_request](#) callback function.

27 **Summary**

28 Provide a function by which the local PMIx server library can return connection and other data
29 posted by local application processes to the host resource manager.

1 **Format**

PMIx v1.0

C

```
2       typedef void (*pmix_dmodex_response_fn_t)(pmix_status_t status,  
3                                                       char *data, size_t sz,  
4                                                       void *cbdata);
```

C

- 5 **IN status**
6 Returned status of the request ([pmix_status_t](#))
- 7 **IN data**
8 Pointer to a data "blob" containing the requested information (handle)
- 9 **IN sz**
10 Number of bytes in the *data* blob (integer)
- 11 **IN cbdata**
12 Data passed into the initial call to [PMIx_server_dmodex_request](#) (memory
13 reference)

14 **Description**

15 Define a function to be called by the PMIx server library for return of information posted by a local
16 application process (via [PMIx_Put](#) with subsequent [PMIx_Commit](#)) in response to a request
17 from the host RM. The returned *data* blob is owned by the PMIx server library and will be free'd
18 upon return from the function.

19 **3.5.13 pmix_connection_cbfunc_t**

20 **Summary**

21 Callback function for incoming connection request from a local client

22 **Format**

PMIx v1.0

C

```
23       typedef void (*pmix_connection_cbfunc_t)(  
24                                                       int incoming_sd, void *cbdata)
```

C

- 25 **IN incoming_sd**
26 (integer)
- 27 **IN cbdata**
28 (memory reference)

1 **Description**

2 Callback function for incoming connection requests from local clients - only used by host
3 environments that wish to directly handle socket connection requests.

4 **3.5.14 pmix_tool_connection_cbfunc_t**

5 **Summary**

6 Callback function for incoming tool connections.

7 **Format**

PMIx v2.0

```

▼────────────────────────────────────────── C ───────────────────────────────────────────▼
8 typedef void (*pmix_tool_connection_cbfunc_t) (
9             pmix_status_t status,
10            pmix_proc_t *proc, void *cbdata)
▲────────────────────────────────────────── C ───────────────────────────────────────────▲

```

- 11 **IN status**
12 `pmix_status_t` value (handle)
- 13 **IN proc**
14 `pmix_proc_t` structure containing the identifier assigned to the tool (handle)
- 15 **IN cbdata**
16 Data to be passed (memory reference)

17 **Description**

18 Callback function for incoming tool connections. The host environment shall provide a
19 namespace/rank identifier for the connecting tool.

▼── Advice to PMIx server hosts ───▼

20 It is assumed that **rank=0** will be the normal assignment, but allow for the future possibility of a
21 parallel set of tools connecting, and thus each process requiring a unique rank.







22 **3.5.15 Constant String Functions**

23 Provide a string representation for several types of values. Note that the provided string is statically
24 defined and must NOT be **free**'d.





1 **Summary**
2 String representation of a [pmix_status_t](#) .

PMIx v1.0

▼  C 
3 **const char***
4 **PMIx_Error_string**(pmix_status_t status);
▲  C 





5 **Summary**
6 String representation of a [pmix_proc_state_t](#) .

PMIx v2.0

▼  C 
7 **const char***
8 **PMIx_Proc_state_string**(pmix_proc_state_t state);
▲  C 





9 **Summary**
10 String representation of a [pmix_scope_t](#) .

PMIx v2.0

▼  C 
11 **const char***
12 **PMIx_Scope_string**(pmix_scope_t scope);
▲  C 





13 **Summary**
14 String representation of a [pmix_persistence_t](#) .

PMIx v2.0

▼  C 
15 **const char***
16 **PMIx_Persistence_string**(pmix_persistence_t persist);
▲  C 

17 **Summary**
18 String representation of a [pmix_data_range_t](#) .

PMIx v2.0

▼  C 
19 **const char***
20 **PMIx_Data_range_string**(pmix_data_range_t range);
▲  C 

1 **Summary**

2 String representation of a `pmix_info_directives_t` .

PMIx v2.0 ▼ _____ C _____ ▼

3 **const char***

4 **PMIx_Info_directives_string**(`pmix_info_directives_t directives`);

▲ _____ C _____ ▲

5 **Summary**

6 String representation of a `pmix_data_type_t` .

PMIx v2.0 ▼ _____ C _____ ▼

7 **const char***

8 **PMIx_Data_type_string**(`pmix_data_type_t type`);

▲ _____ C _____ ▲

9 **Summary**

10 String representation of a `pmix_alloc_directive_t` .

PMIx v2.0 ▼ _____ C _____ ▼

11 **const char***

12 **PMIx_Alloc_directive_string**(`pmix_alloc_directive_t directive`);

▲ _____ 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**

```
PMIx v1.0   ▼ _____ C _____ ▼  
8           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.

▼ _____ **Advice to users** _____ ▼

14 The PMIx *ad hoc* v1.0 Standard defined the `PMIx_Init` function, but modified the function
15 signature in the v1.2 version. The *ad hoc* v1.0 version is not included in this document to avoid
16 confusion.

▲ _____ ▲

17 **4.2.1 PMIx_Init**

18 **Summary**

19 Initialize the PMIx client library

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

Format

PMIx v1.2

```

pmix_status_t
PMIx_Init (pmix_proc_t *proc,
           pmix_info_t info[], size_t ninfo)

```

INOUT `proc`

proc 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.

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_SINGLE_LISTENER "`pmix.sing.listnr`" (`bool`)

Use only one rendezvous socket, letting priorities and/or environment parameters select the active transport. If the library supports multiple methods for clients to connect to servers, this attribute may be supported for disabling all but one of them.

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.

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.

¹<http://libevent.org/>



1 Each call to `PMIx_Init` may contain an array of `pmix_info_t` structures passing directives to
2 the PMIx client library as per the above attributes.
3 Multiple calls to `PMIx_Init` shall not include conflicting directives. The `PMIx_Init` function
4 will return an error when directives that conflict with prior directives are encountered.

5 4.2.2 `PMIx_Finalize`

6 Summary

7 Finalize the PMIx client library.

8 Format

PMIx v1.0  

```
9 pmix_status_t  
10 PMIx_Finalize(const pmix_info_t info[], size_t ninfo)
```

11 **IN** `info`
12 Array of `pmix_info_t` structures (array of handles)
13 **IN** `ninfo`
14 Number of element in the `info` array (`size_t`)

15 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Optional Attributes

16 The following attributes are optional for implementers of PMIx libraries:

17 **PMIX_EMBED_BARRIER** "`pmix.embed.barrier`" (`bool`)
18 Execute a blocking fence operation before executing the specified operation. For example,
19 `PMIx_Finalize` does not include an internal barrier operation by default. This attribute
20 would direct `PMIx_Finalize` to execute a barrier as part of the finalize operation.

21 Description

22 Decrement the PMIx client library reference count. When the reference count reaches zero, the
23 library will finalize the PMIx client, closing the connection with the local PMIx server and
24 releasing all internally allocated memory.

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)

Time in seconds between connection attempts to a PMIx server.

PMIX_CONNECT_MAX_RETRIES "pmix.tool.mretries" (uint32_t)

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

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.

1 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)
 2 Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections,
 3 this attribute may be supported for disabling it.

4 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)
 5 Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections,
 6 this attribute may be supported for disabling it.

7 **PMIX_EVENT_BASE** "pmix.evbase" (struct event_base *)
 8 Pointer to libevent² **event_base** to use in place of the internal progress thread.

9 **PMIX_GDS_MODULE** "pmix.gds.mod" (char*)
 10 Comma-delimited string of desired modules. This attribute is specific to the PRI and
 11 controls only the selection of GDS module for internal use by the process. Module selection
 12 for interacting with the server is performed dynamically during the connection process.

▲-----▲

13 Description

14 Initialize the PMIx tool, returning the process identifier assigned to this tool in the provided
 15 **pmix_proc_t** struct. The *info* array is used to pass user requests pertaining to the init and
 16 subsequent operations. Passing a **NULL** value for the array pointer is supported if no directives are
 17 desired.

18 If called with the **PMIX_TOOL_DO_NOT_CONNECT** attribute, the PMIx tool library will fully
 19 initialize but not attempt to connect to a PMIx server. The tool can connect to a server at a later
 20 point in time, if desired. In all other cases, the PMIx tool library will attempt to connect to
 21 according to the following precedence chain:

- 22 ● if **PMIX_SERVER_URI** or **PMIX_TCP_URI** is given, then connection will be attempted to the
 23 server at the specified URI. Note that it is an error for both of these attributes to be specified.
 24 **PMIX_SERVER_URI** is the preferred method as it is more generalized — **PMIX_TCP_URI** is
 25 provided for those cases where the user specifically wants to use a TCP transport for the
 26 connection and wants to error out if it isn't available or cannot succeed. The PMIx library will
 27 return an error if connection fails — it will not proceed to check for other connection options as
 28 the user specified a particular one to use
- 29 ● if **PMIX_SERVER_PIDINFO** was provided, then the tool will search under the directory
 30 provided by the **PMIX_SERVER_TMPDIR** environmental variable for a rendezvous file created
 31 by the process corresponding to that PID. The PMIx library will return an error if the rendezvous
 32 file cannot be found, or the connection is refused by the server

²<http://libevent.org/>

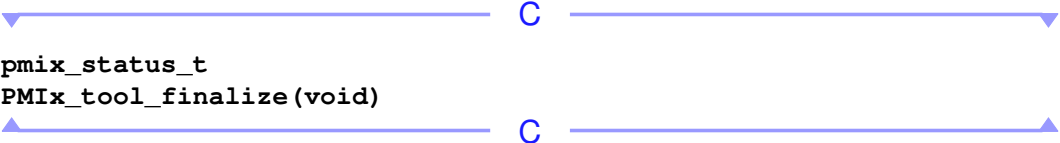
- 1 • if **PMIX_CONNECT_TO_SYSTEM** is given, then the tool will search for a system-level
2 rendezvous file created by a PMIx server in the directory specified by the
3 **PMIX_SYSTEM_TMPDIR** environmental variable. If found, then the tool will attempt to
4 connect to it. An error is returned if the rendezvous file cannot be found or the connection is
5 refused.
- 6 • if **PMIX_CONNECT_SYSTEM_FIRST** is given, then the tool will search for a system-level
7 rendezvous file created by a PMIx server in the directory specified by the
8 **PMIX_SYSTEM_TMPDIR** environmental variable. If found, then the tool will attempt to
9 connect to it. In this case, no error will be returned if the rendezvous file is not found or
10 connection is refused — the PMIx library will silently continue to the next option
- 11 • by default, the tool will search the directory tree under the directory provided by the
12 **PMIX_SERVER_TMPDIR** environmental variable for rendezvous files of PMIx servers,
13 attempting to connect to each it finds until one accepts the connection. If no rendezvous files are
14 found, or all contacted servers refuse connection, then the PMIx library will return an error.
- 15 If successful, the function will return **PMIX_SUCCESS** and will fill the provided structure (if
16 provided) with the server-assigned namespace and rank of the tool. Note that each connection
17 attempt in the above precedence chain will retry (with delay between each retry) a number of times
18 according to the values of the corresponding attributes. Default is no retries.
- 19 Note that the PMIx tool library is referenced counted, and so multiple calls to **PMIx_tool_init**
20 are allowed. Thus, one way to obtain the namespace and rank of the process is to simply call
21 **PMIx_tool_init** with a non-NULL parameter.

22 4.3.2 PMIx_tool_finalize

23 Summary

24 Finalize the PMIx library for a tool connection.

25 Format

26 *PMIx v2.0* 

27 **pmix_status_t**
PMIx_tool_finalize(void)

28 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

29 Description

30 Finalize the PMIx tool library, closing the connection to the server. An error code will be returned
31 if, for some reason, the connection cannot be cleanly terminated — in this case, the connection is
32 dropped.

1 The host RM wants to declare itself as willing to accept tool connection requests.

2 **PMIX_SERVER_SYSTEM_SUPPORT** "pmix.srvr.sys" (bool)

3 The host RM wants to declare itself as being the local system server for PMIx connection
4 requests.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

5 The following attributes are optional for implementers of PMIx libraries:

6 **PMIX_USOCK_DISABLE** "pmix.usock.disable" (bool)

7 Disable legacy UNIX socket (usock) support. If the library supports Unix socket
8 connections, this attribute may be supported for disabling it.

9 **PMIX_SOCKET_MODE** "pmix.sockmode" (uint32_t)

10 POSIX *mode_t* (9 bits valid). If the library supports socket connections, this attribute may
11 be supported for setting the socket mode.

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

13 If provided, directs that the TCP URI be reported and indicates the desired method of
14 reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket
15 connections, this attribute may be supported for reporting the URI.

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

17 Comma-delimited list of devices and/or CIDR notation to include when establishing the
18 TCP connection. If the library supports TCP socket connections, this attribute may be
19 supported for specifying the interfaces to be used.

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

21 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
22 TCP connection. If the library supports TCP socket connections, this attribute may be
23 supported for specifying the interfaces that are *not* to be used.

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

25 The IPv4 port to be used. If the library supports IPV4 connections, this attribute may be
26 supported for specifying the port to be used.

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

28 The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be
29 supported for specifying the port to be used.

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

31 Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections,
32 this attribute may be supported for disabling it.

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

34 Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections,
35 this attribute may be supported for disabling it.

1 **PMIX_SERVER_REMOTE_CONNECTIONS** "pmix.srvr.remote" (bool)
 2 Allow connections from remote tools. Forces the PMIx server to not exclusively use
 3 loopback device. If the library supports connections from remote tools, this attribute may
 4 be supported for enabling or disabling it.

5 **PMIX_EVENT_BASE** "pmix.evbase" (struct event_base *)
 6 Pointer to libevent³ event_base to use in place of the internal progress thread.

7 **PMIX_GDS_MODULE** "pmix.gds.mod" (char*)
 8 Comma-delimited string of desired modules. This attribute is specific to the PRI and
 9 controls only the selection of GDS module for internal use by the process. Module selection
 10 for interacting with the server is performed dynamically during the connection process.

11 Description

12 Initialize the PMIx server support library, and provide a pointer to a `pmix_server_module_t`
 13 structure containing the caller's callback functions. The array of `pmix_info_t` structs is used to
 14 pass additional info that may be required by the server when initializing. For example, it may
 15 include the `PMIX_SERVER_TOOL_SUPPORT` key, thereby indicating that the daemon is willing
 16 to accept connection requests from tools.

Advice to PMIx server hosts

17 Providing a value of `NULL` for the *module* argument is permitted, as is passing an empty *module*
 18 structure. Doing so indicates that the host environment will not provide support for multi-node
 19 operations such as `PMIx_Fence`, but does intend to support local clients access to information.

20 4.4.2 PMIx_server_finalize

21 Summary

22 Finalize the PMIx server library.

23 Format

PMIx v1.0

24 `pmix_status_t`
 25 `PMIx_server_finalize(void)`

26 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

³<http://libevent.org/>

1
2
3

Description

Finalize the PMIx server support library, terminating all connections to attached tools and any local clients. All memory usage is released.

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 publishing/exchanging data must be supported until that objective can be met. This
4 method relies on processes to discover and publish their local information which is collected by the
5 local PMIx server library. Global exchange of the published information is then executed via a
6 collective operation performed by the host SMS servers.

7 Keys are required to be unique within a specific level of information as defined in 3.4.10. For
8 example, a value for `PMIX_NUM_NODES` can be specified for each of the `session`, `job`, and
9 `application` levels. However, subsequently specifying another value for that attribute in the
10 `session` level will overwrite the prior value.

5.1 Setting and Accessing Key/Value Pairs

5.1.1 PMIx_Put

Summary

14 Push a key/value pair into the client's namespace.

Format

15 *PMIx v1.0*

```
16 pmix_status_t  
17 PMIx_Put (pmix_scope_t scope,  
18           const pmix_key_t key,  
19           pmix_value_t *val)
```

20 **IN** `scope`
21 Distribution scope of the provided value (handle)

22 **IN** `key`
23 key (`pmix_key_t`)

24 **IN** `value`
25 Reference to a `pmix_value_t` structure (handle)

26 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 27. 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

Format

PMIx v1.0

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 the PMIx server if not found.

20

21

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

22

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.

23

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 one containing the requesting process is desired, then the attribute array must contain a

28

`PMIX_SESSION_ID` attribute identifying the desired target.

29

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

Format

PMIx v1.0

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

7

process reference (handle)

8

IN key

9

key to retrieve (string)

10

IN info

11

Array of info structures (array of handles)

12

IN ninfo

13

Number of elements in the *info* array (integer)

14

IN cbfunc

15

Callback function (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 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning from the API.

21

22

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

23

24

- 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

25

26

If executed, the status returned in the provided callback function will be one of the following constants:

27

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)

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

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

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

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

7 Scope of the data to be found in a **PMIx_Get** call.

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

9 Return information about the specified session. If information about a session other than the
10 one containing the requesting process is desired, then the attribute array must contain a
11 **PMIX_SESSION_ID** attribute identifying the desired target.

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

13 Return information about the specified job or namespace. If information about a job or
14 namespace other than the one containing the requesting process is desired, then the attribute
15 array must contain a **PMIX_JOBID** or **PMIX_NAMESPACE** attribute identifying the desired
16 target. Similarly, if information is requested about a job or namespace in a session other than
17 the one containing the requesting process, then an attribute identifying the target session
18 must be provided.

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

20 Return information about the specified application. If information about an application other
21 than the one containing the requesting process is desired, then the attribute array must
22 contain a **PMIX_APPNUM** attribute identifying the desired target. Similarly, if information is
23 requested about an application in a job or session other than the one containing the requesting
24 process, then attributes identifying the target job and/or session must be provided.

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

26 Return information about the specified node. If information about a node other than the one
27 containing the requesting process is desired, then the attribute array must contain either the
28 **PMIX_NODEID** or **PMIX_HOSTNAME** attribute identifying the desired target.

▲-----▲
▼-----▼ **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 delivery of the data by the host environment 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

The callback function will be executed once the specified data becomes available from the identified process and retrieved by the local server. The *info* array is used as described by the `PMIx_Get` routine.

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`).

In general, only 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`. See 3.4.10 for an explanation regarding use of the *level* attributes.

5.1.4 PMIx_Store_internal

Summary

Store some data locally for retrieval by other areas of the proc.

1 **Format**

PMIx v1.0

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**
process reference (handle)
- 8 **IN** **key**
key to retrieve (string)
- 10 **IN** **val**
Value to store (handle)

12 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

13 **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.nspace, 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;

- for non-specific attributes (e.g., `PMIX_NUM_NODES`), including the `PMIX_APP_INFO` attribute to indicate that the application-level information for that attribute is being requested

Example requests are shown below:

```

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
36
37

```

```

C

```

```

pmix_info_t info;
pmix_value_t *value;
pmix_status_t rc;
pmix_proc_t myproc, otherproc;
uint32_t appsize, appnum;

/* initialize the client library */
PMIx_Init(&myproc, NULL, 0);

/* get the #processes in our application */
rc = PMIx_Get(&myproc, PMIX_APP_SIZE, NULL, 0, &value);
appsize = value->data.uint32;

/* get the #nodes in an application containing "otherproc".
 * Note that the rank of a process in the other application
 * must be obtained first - a simple method is shown here */

/* assume for this example that we are in the first application
 * and we want the #nodes in the second application - use the
 * rank of the first process in that application, remembering
 * that ranks start at zero */
PMIX_PROC_LOAD(&otherproc, myproc.nspace, appsize);

PMIX_INFO_LOAD(&info, PMIX_APP_INFO, NULL, PMIX_BOOL);
rc = PMIx_Get(&otherproc, PMIX_NUM_NODES, &info, 1, &value);

/* alternatively, we can directly ask for the #nodes in
 * the second application in our job, again remembering that
 * application numbers start with zero */
appnum = 1;
PMIX_INFO_LOAD(&appinfo[0], PMIX_APP_INFO, NULL, PMIX_BOOL);
PMIX_INFO_LOAD(&appinfo[1], PMIX_APPNUM, &appnum, PMIX_UINT32);
rc = PMIx_Get(&myproc, PMIX_NUM_NODES, appinfo, 2, &value);

```

```

C

```

1 5.1.5.4 Process-level information

2 Process-level information is accessed by providing the namespace and rank of the target process. In
3 the absence of any directive as to the level of information being requested, the PMIx library will
4 always return the process-level value.

5 5.1.5.5 Node-level information

6 Information regarding a node within the system can be obtained by the following methods:

- 7 • for node-level attributes (e.g., [PMIX_NODE_SIZE](#)), specifying the namespace and rank of a
8 process executing on the target node;
- 9 • for node-level attributes (e.g., [PMIX_NODE_SIZE](#)), including the [PMIX_NODEID](#) or
10 [PMIX_HOSTNAME](#) attribute specifying the node whose information is being requested. In this
11 case, the *proc* argument's values are ignored; or
- 12 • for non-specific attributes (e.g., [PMIX_NUM_SLOTS](#)), including the [PMIX_NODE_INFO](#)
13 attribute to indicate that the node-level information for that attribute is being requested

14 Example requests are shown below:

```
15 pmix_info_t info[2];  
16 pmix_value_t *value;  
17 pmix_status_t rc;  
18 pmix_proc_t myproc, otherproc;  
19 uint32_t nodeid;  
20  
21 /* initialize the client library */  
22 PMIx_Init(&myproc, NULL, 0);  
23  
24 /* get the #procs on our node */  
25 rc = PMIx_Get(&myproc, PMIX_NODE_SIZE, NULL, 0, &value);  
26  
27 /* get the #slots on another node */  
28 PMIX_INFO_LOAD(&info[0], PMIX_NODE_INFO, NULL, PMIX_BOOL);  
29 PMIX_INFO_LOAD(&info[1], PMIX_HOSTNAME, "remotehost", PMIX_STRING);  
30 rc = PMIx_Get(&myproc, PMIX_NUM_SLOTS, info, 2, &value);  
31
```

Advice to users

32 An explanation of the use of [PMIx_Get](#) versus [PMIx_Query_info_nb](#) is provided in [7.1.3.1](#).

1 5.2 Exchanging Key/Value Pairs

2 The APIs defined in this section push key/value pairs from the client to the local PMIx server, and
3 circulate the data between PMIx servers for subsequent retrieval by the local clients.



4 5.2.1 PMIx_Commit

5 Summary

6 Push all previously **PMIx_Put** values to the local PMIx server.

7 Format

8 *PMIx v1.0*

▼  ▼
`pmix_status_t PMIx_Commit(void)`
▲  ▲

9 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

10 Description

11 This is an asynchronous operation. The PRI will immediately return to the caller while the data is
12 transmitted to the local server in the background.

▼ **Advice to users** ▼

13 The local PMIx server will cache the information locally - i.e., the committed data will not be
14 circulated during **PMIx_Commit** . Availability of the data upon completion of **PMIx_Commit** is
15 therefore implementation-dependent.

16 5.2.2 PMIx_Fence

17 Summary

18 Execute a blocking barrier across the processes identified in the specified array, collecting
19 information posted via **PMIx_Put** as directed.

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

Format

PMIx v1.0

```

pmix_status_t
PMIx_Fence(const pmix_proc_t procs[], size_t nprocs,
           const pmix_info_t info[], size_t ninfo)

```

- IN procs**
Array of `pmix_proc_t` structures (array of handles)
- IN nprocs**
Number of element in the `procs` array (integer)
- IN info**
Array of info structures (array of handles)
- IN ninfo**
Number of element in the `info` 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_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 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.

10 The *info* array is used to pass user requests regarding the fence operation.

11 Note that for scalability reasons, the default behavior for `PMIx_Fence` is to *not* collect the data.

Advice to PMIx library implementers

12 `PMIx_Fence` and its non-blocking form are both *collective* operations. Accordingly, the PMIx
13 server library is required to aggregate participation by local clients, passing the request to the host
14 environment once all local participants have executed the API.

Advice to PMIx server hosts

15 The host will receive a single call for each collective operation. It is the responsibility of the host to
16 identify the nodes containing participating processes, execute the collective across all participating
17 nodes, and notify the local PMIx server library upon completion of the global collective.

19 5.2.3 `PMIx_Fence_nb`

20 Summary

21 Execute a nonblocking `PMIx_Fence` across the processes identified in the specified array of
22 processes, collecting information posted via `PMIx_Put` as directed.

1

Format

PMIx v1.0

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

7

Array of **pmix_proc_t** structures (array of handles)

8

IN nprocs

9

Number of element 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 (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 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning from the API.

21

22

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called. This can occur if the collective involved only processes on the local node.

23

24

25

- 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

26

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.

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](#) .

1

Format

PMIx v1.0

2

`pmix_status_t`

3

`PMIx_Publish(const pmix_info_t info[], size_t ninfo)`

4

IN `info`

5

Array of info structures (array of handles)

6

IN `ninfo`

7

Number of element in the *info* array (integer)

8

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Required Attributes

9

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 published the info.

10

11

12

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 the target process from ever exposing its data.

17

18

PMIX_RANGE "`pmix.range`" (`pmix_data_range_t`)

19

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

20

PMIX_PERSISTENCE "`pmix.persist`" (`pmix_persistence_t`)

21

Value for calls to `PMIx_Publish`.

Advice to PMIx library implementers

22

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.

23

24

25

26

27

Description

Publish the data in the *info* array for subsequent lookup. By default, the data will be published into the `PMIX_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

Duplicate keys within the specified data range may lead to unexpected behavior depending on host RM implementation of the backing key-value store.

Advice to PMIx library implementers

Implementations should, to the best of their ability, detect duplicate keys and protect the user from unexpected behavior - preferably returning an error. This version of the standard does not define a specific error code to be returned, so the implementation must make it clear to the user what to expect in this scenario. One suggestion is to define an RM specific error code beyond the `PMIX_EXTERNAL_ERR_BASE` boundary. Future versions of the standard will clarify that a specific PMIx error be returned when conflicting values are published for a given key, and will provide attributes to allow modified behaviors such as overwrite.

5.3.2 `PMIx_Publish_nb`

Summary

Nonblocking `PMIx_Publish` routine.

1

Format

PMIx v1.0

C

2

`pmix_status_t`

3

`PMIx_Publish_nb(const pmix_info_t info[], size_t ninfo,`

4

`pmix_op_cbfunc_t cbfunc, void *cbdata)`

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

IN cbfunc

10

Callback function `pmix_op_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 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning from the API.

15

16

17

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

18

19

- 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

20

Required Attributes

21

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 published the info.

22

23

24

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
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

Format

PMIx v1.0

```

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_SESSION` range. Changes to the range, and any additional directives, can be provided in the `pmix_info_t` array.

Note that the search is also constrained to only data published by the current user (i.e., the search will not return data published by an application being executed by another user). 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.

5.3.4 `PMIx_Lookup_nb`

Summary

Nonblocking version of `PMIx_Lookup`.

1

Format

PMIx v1.0

C

2

```

pmix_status_t
PMIx_Lookup_nb(char **keys,
               const pmix_info_t info[], size_t ninfo,
               pmix_lookup_cbfunc_t cbfunc, void *cbdata)

```

5

C

6

IN keys

7

Array to be provided to the callback (array of strings)

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 (handle)

14

IN cbdata

15

Callback data to be provided to the callback function (pointer)

16

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 library *must not* invoke the callback function prior to returning from the API.

18

19

- a PMIx error constant indicating an error in the input - the *cbfunc* will *not* be called

20

Required Attributes

21

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.

22

23

24

Optional Attributes

25

The following attributes are optional for host environments that support this operation:

26

PMIX_TIMEOUT "pmix.timeout" (int)

27

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.

28

29

30

PMIX_RANGE "pmix.range" (pmix_data_range_t)

31

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

32

PMIX_WAIT "pmix.wait" (int)

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

▲-----▲
▼-----▼ **Advice to PMIx library implementers** -----▼

3 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
4 environment due to race condition considerations between completion of the operation versus
5 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
6 directly in the PMIx server library must take care to resolve the race condition and should avoid
7 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
8 created.

9 **Description**

10 Non-blocking form of the **PMIx_Lookup** function. Data for the provided NULL-terminated *keys*
11 array will be returned in the provided callback function. As with **PMIx_Lookup**, the default
12 behavior is to *not* wait for data to be published. The *info* array can be used to modify the behavior
13 as previously described by **PMIx_Lookup**. Both the *info* and *keys* arrays must be maintained until
14 the callback is provided.

15 **5.3.5 PMIx_Unpublish**

16 **Summary**

17 Unpublish data posted by this process using the given keys.

1

Format

PMIx v1.0

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 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.

13

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 the target process from ever exposing its data.

18

19

PMIX_RANGE "`pmix.range`" (`pmix_data_range_t`)

20

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

21

Advice to PMIx library implementers

22

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.

26

Description

Unpublish data posted by this process using the given *keys*. The function will block until the data has been removed by the server (i.e., it is safe to publish that key again). A value of **NULL** for the *keys* parameter instructs the server to remove *all* data published by this process.

By default, the range is assumed to be **PMIX_SESSION**. Changes to the range, and any additional directives, can be provided in the *info* array.

5.3.6 PMIx_Unpublish_nb

Summary

Nonblocking version of **PMIx_Unpublish**.

Format

PMIx v1.0

```
pmix_status_t
PMIx_Unpublish_nb(char **keys,
                  const pmix_info_t info[], size_t ninfo,
                  pmix_op_cbfunc_t cbfunc, void *cbdata)
```

IN keys
(array of strings)

IN info
Array of info structures (array of handles)

IN ninfo
Number of element 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 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

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.

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 **Format**

PMIx v1.0

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`
Array of info structures (array of handles)
- 7
- 8 **IN** `ninfo`
Number of elements in the `job_info` array (integer)
- 9
- 10 **IN** `apps`
Array of `pmix_app_t` structures (array of handles)
- 11
- 12 **IN** `napps`
Number of elements in the `apps` array (integer)
- 13
- 14 **OUT** `nspace`
Namespace of the new job (string)
- 15

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 Host environments that implement support for `PMIx_Spawn` are required to pass the
29 `PMIX_SPAWNED` and `PMIX_PARENT_ID` attributes to all PMIx servers launching new child
30 processes so those values can be returned to clients upon connection to the PMIx server. In
31 addition, they are required to support the following attributes when present in either the `job_info` or
32 the `info` array of an element of the `apps` array:

33 **PMIX_WDIR** "pmix.wdir" (char*)
34 Working directory for spawned processes.

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

5 **PMIX_PREFIX** "pmix.prefix" (char*)
6 Prefix to use for starting spawned processes.

7 **PMIX_HOST** "pmix.host" (char*)
8 Comma-delimited list of hosts to use for spawned processes.

9 **PMIX_HOSTFILE** "pmix.hostfile" (char*)
10 Hostfile to use for spawned processes.

▲-----▲

▼-----▼ **Optional Attributes** -----▼

11 The following attributes are optional for host environments that support this operation:

12 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)
13 Hostfile listing hosts to add to existing allocation.

14 **PMIX_ADD_HOST** "pmix.addhost" (char*)
15 Comma-delimited list of hosts to add to the allocation.

16 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)
17 Preload binaries onto nodes.

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

20 **PMIX_PERSONALITY** "pmix.pers" (char*)
21 Name of personality to use.

22 **PMIX_MAPPER** "pmix.mapper" (char*)
23 Mapping mechanism to use for placing spawned processes - when accessed using
24 **PMIx_Get** , use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping
25 mechanism used for the provided namespace.

26 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)
27 Display process mapping upon spawn.

28 **PMIX_PPR** "pmix.ppr" (char*)
29 Number of processes to spawn on each identified resource.

30 **PMIX_MAPBY** "pmix.mapby" (char*)
31 Process mapping policy - when accessed using **PMIx_Get** , use the
32 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
33 provided namespace

34 **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 for the
3 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 the
7 provided namespace

8 **PMIX_NON_PMI** "pmix.nonpmi" (bool)
9 Spawned processes will not call `PMIx_Init` .

10 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)
11 Spawned process rank that is to receive `stdin`.

12 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (bool)
13 Forward this process's `stdin` to the designated process.

14 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)
15 Forward `stdout` from spawned processes to this process.

16 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)
17 Forward `stderr` from spawned processes to this process.

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

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

22 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
23 Timestamp output from applications.

24 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)
25 Merge `stdout` and `stderr` streams from application processes.

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

28 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)
29 Mark the `argv` with the rank of the process.

30 **PMIX_CPUS_PER_PROC** "pmix.cpusperproc" (uint32_t)
31 Number of cpus to assign to each rank - when accessed using `PMIx_Get` , use the
32 `PMIX_RANK_WILDCARD` value for the rank to discover the cpus/process assigned to the
33 provided namespace

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

36 **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 Description

17 Spawn a new job. The assigned namespace of the spawned applications is returned in the *nspace*
18 parameter. A **NULL** value in that location indicates that the caller doesn't wish to have the
19 namespace returned. The *nspace* array must be at least of size one more than **PMIX_MAX_NSLEN** .

20 By default, the spawned processes will be PMIx "connected" to the parent process upon successful
21 launch (see **PMIx_Connect** description for details). Note that this only means that (a) the parent
22 process will be given a copy of the new job's information so it can query job-level info without
23 incurring any communication penalties, (b) newly spawned child processes will receive a copy of
24 the parent processes job-level info, and (c) both the parent process and members of the child job
25 will receive notification of errors from processes in their combined assemblage.

▼ Advice to users ▼

26 Behavior of individual resource managers may differ, but it is expected that failure of any
27 application process to start will result in termination/cleanup of *all* processes in the newly spawned
28 job and return of an error code to the caller.



1 6.2.2 PMIx_Spawn_nb

2 Summary

3 Nonblocking version of the `PMIx_Spawn` routine.

4 Format

PMIx v1.0

C

```
5 pmix_status_t
6 PMIx_Spawn_nb(const pmix_info_t job_info[], size_t ninfo,
7               const pmix_app_t apps[], size_t napps,
8               pmix_spawn_cbfunc_t cbfunc, void *cbdata)
```

C

- 9 **IN** `job_info`
10 Array of info structures (array of handles)
- 11 **IN** `ninfo`
12 Number of elements in the `job_info` array (integer)
- 13 **IN** `apps`
14 Array of `pmix_app_t` structures (array of handles)
- 15 **IN** `cbfunc`
16 Callback function `pmix_spawn_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 library *must not* invoke the callback
22 function prior to returning from the API.
- 23 • a PMIx error constant indicating an error in the request - 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 following attributes to those provided before passing the request to the host:

- 27 **PMIX_SPAWNED** "`pmix.spawned`" (bool)
28 **true** if this process resulted from a call to `PMIx_Spawn`.
- 29 **PMIX_PARENT_ID** "`pmix.parent`" (`pmix_proc_t`)
30 Process identifier of the parent process of the calling process.
- 31 **PMIX_REQUESTOR_IS_CLIENT** "`pmix.req.client`" (bool)
32 The requesting process is a PMIx client.

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

2 The requesting process is a PMIx tool.

3 Host environments that implement support for **PMIx_Spawn** are required to pass the
4 **PMIX_SPAWNED** and **PMIX_PARENT_ID** attributes to all PMIx servers launching new child
5 processes so those values can be returned to clients upon connection to the PMIx server. In
6 addition, they are required to support the following attributes when present in either the *job_info* or
7 the *info* array of an element of the *apps* array:

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

9 Working directory for spawned processes.

10 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

11 Set the application's current working directory to the session working directory assigned by
12 the RM - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for
13 the rank to discover the session working directory assigned to the provided namespace

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

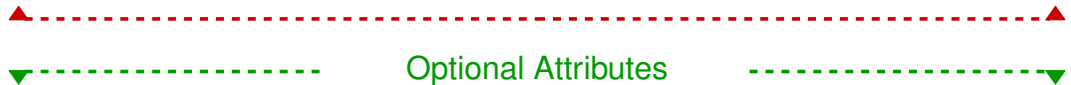
15 Prefix to use for starting spawned processes.

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

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

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

19 Hostfile to use for spawned processes.



20 The following attributes are optional for host environments that support this operation:

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

22 Hostfile listing hosts to add to existing allocation.

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

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

25 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

26 Preload binaries onto nodes.

27 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

28 Comma-delimited list of files to pre-position on nodes.

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

30 Name of personality to use.

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

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

1 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)
2 Display process mapping upon spawn.

3 **PMIX_PPR** "pmix.ppr" (char*)
4 Number of processes to spawn on each identified resource.

5 **PMIX_MAPBY** "pmix.mapby" (char*)
6 Process mapping policy - when accessed using **PMIx_Get** , use the
7 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
8 provided namespace

9 **PMIX_RANKBY** "pmix.rankby" (char*)
10 Process ranking policy - when accessed using **PMIx_Get** , use the
11 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the
12 provided namespace

13 **PMIX_BINDTO** "pmix.bindto" (char*)
14 Process binding policy - when accessed using **PMIx_Get** , use the
15 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
16 provided namespace

17 **PMIX_NON_PMI** "pmix.nonpmi" (bool)
18 Spawned processes will not call **PMIx_Init** .

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

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

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

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

27 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)
28 Spawned application consists of debugger daemons.

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

31 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
32 Timestamp output from applications.

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

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

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

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

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

9 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)
10 Do not oversubscribe the cpus.

11 **PMIX_REPORT_BINDINGS** "pmix.repbinding" (bool)
12 Report bindings of the individual processes.

13 **PMIX_CPU_LIST** "pmix.cpulist" (char*)
14 List of cpus to use for this job - when accessed using **PMIx_Get** , use the
15 **PMIX_RANK_WILDCARD** value for the rank to discover the cpu list used for the provided
16 namespace

17 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)
18 Application supports recoverable operations.

19 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)
20 Application is continuous, all failed processes should be immediately restarted.

21 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)
22 Maximum number of times to restart a job - when accessed using **PMIx_Get** , use the
23 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
24 namespace



25 Description

26 Nonblocking version of the **PMIx_Spawn** routine. The provided callback function will be
27 executed upon successful start of *all* specified application processes.



Advice to users

28 Behavior of individual resource managers may differ, but it is expected that failure of any
29 application process to start will result in termination/cleanup of *all* processes in the newly spawned
30 job and return of an error code to the caller.



1 6.3 Connecting and Disconnecting Processes

2 This section defines functions to connect and disconnect processes in two or more separate PMIx
3 namespaces. The PMIx definition of *connected* solely implies the following:

- 4 • job-level information for each namespace involved in the operation is to be made available to all
5 processes in the connected assemblage
- 6 • any data posted by a process in the connected assemblage (via calls to **PMIx_Put** committed via
7 **PMIx_Commit**) prior to execution of the **PMIx_Connect** operation is to be made accessible
8 to all processes in the assemblage - any data posted after execution of the *connect* operation must
9 be exchanged via a separate **PMIx_Fence** operation spanning the connected processes
- 10 • the host environment should treat the failure of any process in the assemblage as a reportable
11 event, taking action on the assemblage as if it were a single application. For example, if the
12 environment defaults (in the absence of any application directives) to terminating an application
13 upon failure of any process in that application, then the environment should terminate all
14 processes in the connected assemblage upon failure of any member.

Advice to PMIx server hosts

15 The host environment may choose to assign a new namespace to the connected assemblage and/or
16 assign new ranks for its members for its own internal tracking purposes. However, it is not required
17 to communicate such assignments to the participants (e.g., in response to an appropriate call to
18 **PMIx_Query_info_nb**). The host environment is required to generate a
19 **PMIX_ERR_INVALID_TERMINATION** event should any process in the assemblage terminate or
20 call **PMIx_Finalize** without first *disconnecting* from the assemblage.

Advice to users

21 Attempting to *connect* processes solely within the same namespace is essentially a *no-op* operation.
22 While not explicitly prohibited, users are advised that a PMIx implementation or host environment
23 may return an error in such cases.

24 Neither the PMIx implementation nor host environment are required to provide any tracking
25 support for the assemblage. Thus, the application is responsible for maintaining the membership
26 list of the assemblage.

27 6.3.1 PMIx_Connect

28 Summary

29 Connect namespaces.

1 **Format**

PMIx v1.0

```

2 pmix_status_t
3 PMIx_Connect(const pmix_proc_t procs[], size_t nprocs,
4             const pmix_info_t info[], size_t ninfo)

```

- 5 **IN procs**
Array of proc structures (array of handles)
- 6
- 7 **IN nprocs**
Number of elements in the *procs* array (integer)
- 8
- 9 **IN info**
Array of info structures (array of handles)
- 10
- 11 **IN ninfo**
Number of elements in the *info* array (integer)
- 12

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.
- 21 **PMIX_COLLECTIVE_ALGO** "pmix.calgo" (char*)
22 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
23 impose any requirements on a host environment’s collective algorithms. Thus, the
24 acceptable values for this attribute will be environment-dependent - users are encouraged to
25 check their host environment for supported values.
- 26 **PMIX_COLLECTIVE_ALGO_REQD** "pmix.calreqd" (bool)
27 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 Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The
8 function will return once all processes identified in *procs* have called either `PMIx_Connect` or
9 its non-blocking version, *and* the host environment has completed any supporting operations
10 required to meet the terms of the PMIx definition of *connected* processes.
11

Advice to users

12 All processes engaged in a given `PMIx_Connect` operation must provide the identical *procs*
13 array as ordering of entries in the array and the method by which those processes are identified
14 (e.g., use of `PMIX_RANK_WILDCARD` versus listing the individual processes) *may* impact the
15 host environment's algorithm for uniquely identifying an operation.

Advice to PMIx library implementers

16 `PMIx_Connect` and its non-blocking form are both *collective* operations. Accordingly, the PMIx
17 server library is required to aggregate participation by local clients, passing the request to the host
18 environment once all local participants have executed the API.

Advice to PMIx server hosts

19 The host will receive a single call for each collective operation. It is the responsibility of the host to
20 identify the nodes containing participating processes, execute the collective across all participating
21 nodes, and notify the local PMIx server library upon completion of the global collective.

1 Processes that combine via **PMIx_Connect** must call **PMIx_Disconnect** prior to finalizing
2 and/or terminating - any process in the assemblage failing to meet this requirement will cause a
3 **PMIX_ERR_INVALID_TERMINATION** event to be generated.

4 A process can only engage in *one* connect operation involving the identical *procs* array at a time.
5 However, a process *can* be simultaneously engaged in multiple connect operations, each involving a
6 different *procs* array.

7 As in the case of the **PMIx_Fence** operation, the *info* array can be used to pass user-level
8 directives regarding the algorithm to be used for any collective operation involved in the operation,
9 timeout constraints, and other options available from the host RM.

10 6.3.2 PMIx_Connect_nb

11 Summary

12 Nonblocking **PMIx_Connect_nb** routine.

13 Format

14 *PMIx v1.0*

C

```
14 pmix_status_t  
15 PMIx_Connect_nb(const pmix_proc_t procs[], size_t nprocs,  
16                 const pmix_info_t info[], size_t ninfo,  
17                 pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

18 **IN procs**
19 Array of proc structures (array of handles)
20 **IN nprocs**
21 Number of elements in the *procs* array (integer)
22 **IN info**
23 Array of info structures (array of handles)
24 **IN ninfo**
25 Number of element in the *info* array (integer)
26 **IN cbfunc**
27 Callback function **pmix_op_cbfunc_t** (function reference)
28 **IN cbdata**
29 Data to be passed to the callback function (memory reference)

30 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 PMIx libraries are not required to directly support any attributes for this function. However, any
9 provided attributes must be passed to the host SMS daemon for processing.

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.

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

16 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
17 impose any requirements on a host environment’s collective algorithms. Thus, the
18 acceptable values for this attribute will be environment-dependent - users are encouraged to
19 check their host environment for supported values.

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

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

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 completion of the operation versus
24 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
25 directly in the PMIx server library must take care to resolve the race condition and should avoid
26 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
27 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.

Format

PMIx v1.0

```
pmix_status_t
PMIx_Disconnect(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 element 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.

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

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

1 The host will receive a single call for each collective operation. The host will receive a single call
2 for each collective operation. It is the responsibility of the host to identify the nodes containing
3 participating processes, execute the collective across all participating nodes, and notify the local
4 PMIx server library upon completion of the global collective.

5 A process can only engage in *one* disconnect operation involving the identical *procs* array at a time.
6 However, a process *can* be simultaneously engaged in multiple disconnect operations, each
7 involving a different *procs* array.

8 As in the case of the **PMIx_Fence** operation, the *info* array can be used to pass user-level
9 directives regarding the algorithm to be used for any collective operation involved in the operation,
10 timeout constraints, and other options available from the host RM.

6.3.4 PMIx_Disconnect_nb

Summary

13 Nonblocking **PMIx_Disconnect** routine.

Format

14 *PMIx v1.0*

```
15 pmix_status_t  
16 PMIx_Disconnect_nb(const pmix_proc_t procs[], size_t nprocs,  
17                   const pmix_info_t info[], size_t ninfo,  
18                   pmix_op_cbfunc_t cbfunc, void *cbdata);
```

19 **IN** **procs**
20 Array of proc structures (array of handles)
21 **IN** **nprocs**
22 Number of elements in the *procs* array (integer)
23 **IN** **info**
24 Array of info structures (array of handles)
25 **IN** **ninfo**
26 Number of element in the *info* array (integer)
27 **IN** **cbfunc**
28 Callback function **pmix_op_cbfunc_t** (function reference)
29 **IN** **cbdata**
30 Data to be passed to the callback function (memory reference)

1 Returns one of the following:

- 2 • **PMIX_SUCCESS** , indicating that the request is being processed by the host environment - result
3 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback
4 function prior to returning from the API.
- 5 • **PMIX_OPERATION_SUCCEEDED** , indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will *not* be called
- 7 • a PMIx error constant indicating either an error in the input or that the request was immediately
8 processed and failed - the *cbfunc* will *not* be called

Required Attributes

9 PMIx libraries are not required to directly support any attributes for this function. However, any
10 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

11 The following attributes are optional for host environments that support this operation:

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

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

Advice to PMIx library implementers

16 We recommend that implementation of the **PMIX_TIMEOUT** attribute be left to the host
17 environment due to race condition considerations between completion of the operation versus
18 internal timeout in the PMIx server library. Implementers that choose to support **PMIX_TIMEOUT**
19 directly in the PMIx server library must take care to resolve the race condition and should avoid
20 passing **PMIX_TIMEOUT** to the host environment so that multiple competing timeouts are not
21 created.

Description

22 Nonblocking **PMIx_Disconnect** routine. The callback function is called once all processes
23 identified in *procs* have called either **PMIx_Disconnect_nb** or its blocking version, *and* the
24 host environment has completed any required supporting operations. See the advice provided in the
25 description for **PMIx_Disconnect** for more information.
26

CHAPTER 7

Job Allocation 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.

1 7.1.1 PMIx_Resolve_peers

2 Summary

3 Obtain the array of processes within the specified namespace that are executing on a given node.

4 Format

PMIx v1.0

C

```
5 pmix_status_t
6 PMIx_Resolve_peers(const char *nodename,
7                   const pmix_namespace_t nspace,
8                   pmix_proc_t **procs, size_t *nprocs)
```

C

9 **IN** **nodename**

10 Name of the node to query (string)

11 **IN** **nspace**

12 namespace (string)

13 **OUT** **procs**

14 Array of process structures (array of handles)

15 **OUT** **nprocs**

16 Number of elements in the *procs* array (integer)

17 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

18 Description

19 Given a *nodename*, return the array of processes within the specified *nspace* that are executing on
20 that node. If the *nspace* is **NULL**, then all processes on the node will be returned. If the specified
21 node does not currently host any processes, then the returned array will be **NULL**, and *nprocs* will
22 be 0. The caller is responsible for releasing the *procs* array when done with it. The
23 **PMIX_PROC_FREE** macro is provided for this purpose.

24 7.1.2 PMIx_Resolve_nodes

25 Summary

26 Return a list of nodes hosting processes within the given namespace.

1 **Format**

PMIx v1.0

C

2 `pmix_status_t`

3 `PMIx_Resolve_nodes(const char *nspace, char **nodelist)`

C

4 **IN** `nspace`

Namespace (string)

6 **OUT** `nodelist`

Comma-delimited list of nodenames (string)

8 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

9 **Description**

10 Given a *nspace*, return the list of nodes hosting processes within that namespace. The returned
11 string will contain a comma-delimited list of nodenames. The caller is responsible for releasing the
12 string when done with it.

13 **7.1.3 PMIx_Query_info_nb**

14 **Summary**

15 Query information about the system in general.

16 **Format**

PMIx v2.0

C

17 `pmix_status_t`

18 `PMIx_Query_info_nb(pmix_query_t queries[], size_t nqueries,`
19 `pmix_info_cbfunc_t cbfunc, void *cbdata)`

C

20 **IN** `queries`

Array of query structures (array of handles)

22 **IN** `nqueries`

Number of elements in the *queries* array (integer)

24 **IN** `cbfunc`

Callback function `pmix_info_cbfunc_t` (function reference)

26 **IN** `cbdata`

Data to be passed to the callback function (memory reference)

28 Returns one of the following:

- **PMIX_SUCCESS** indicating that the request has been accepted for processing and the provided callback function will be executed upon completion of the operation. Note that the library *must not* invoke the callback function prior to returning from the API.
- a non-zero PMIx error constant indicating a reason for the request to have been rejected. In this case, the provided callback function will *not* be executed

If executed, the status returned in the provided callback function will be one of the following constants:

- **PMIX_SUCCESS** All data has been returned
- **PMIX_ERR_NOT_FOUND** None of the requested data was available
- **PMIX_ERR_PARTIAL_SUCCESS** Some of the data has been returned
- **PMIX_ERR_NOT_SUPPORTED** The host RM does not support this function
- a non-zero PMIx error constant indicating a reason for the request's failure

▼----- Required Attributes -----▼

PMIx libraries that support this API are required to support the following attributes:

PMIX_QUERY_REFRESH_CACHE "pmix.qry.rfsh" (bool)

Retrieve updated information from server.

PMIX_SESSION_INFO "pmix.ssn.info" (bool)

Return information about the specified session. If information about a session other than the one containing the requesting process is desired, then the attribute array must contain a **PMIX_SESSION_ID** attribute identifying the desired target.

PMIX_JOB_INFO "pmix.job.info" (bool)

Return information about the specified job or namespace. If information about a job or namespace other than the one containing the requesting process is desired, then the attribute array must contain a **PMIX_JOBID** or **PMIX_NAMESPACE** attribute identifying the desired target. Similarly, if information is requested about a job or namespace in a session other than the one containing the requesting process, then an attribute identifying the target session must be provided.

PMIX_APP_INFO "pmix.app.info" (bool)

Return information about the specified application. If information about an application other than the one containing the requesting process is desired, then the attribute array must contain a **PMIX_APPNUM** attribute identifying the desired target. Similarly, if information is requested about an application in a job or session other than the one containing the requesting process, then attributes identifying the target job and/or session must be provided.

PMIX_NODE_INFO "pmix.node.info" (bool)

1 Return information about the specified node. If information about a node other than the one
2 containing the requesting process is desired, then the attribute array must contain either the
3 **PMIX_NODEID** or **PMIX_HOSTNAME** attribute identifying the desired target.

4 **PMIX_PROCID** "pmix.procid" (**pmix_proc_t**)

5 Process identifier Specifies the process ID whose information is being requested - e.g., a
6 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
7 request is for information on a specific process.

8 **PMIX_NAMESPACE** "pmix.namespace" (**char***)

9 Namespace of the job. Specifies the namespace of the process whose information is being
10 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
11 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
12 information on a specific process.

13 **PMIX_RANK** "pmix.rank" (**pmix_rank_t**)

14 Process rank within the job. Specifies the rank of the process whose information is being
15 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
16 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
17 information on a specific process.

18 Note that inclusion of the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
19 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
20 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on
21 multiple specific processes therefore requires submitting multiple **pmix_query_t** structures,
22 each referencing one process.

23 PMIx libraries are not required to directly support any other attributes for this function. However,
24 any provided attributes must be passed to the host SMS daemon for processing, and the PMIx
25 library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client
26 process making the request.



27 Host environments that support this operation are required to support the following attributes as
28 qualifiers to the request:

29 **PMIX_PROCID** "pmix.procid" (**pmix_proc_t**)

30 Process identifier Specifies the process ID whose information is being requested - e.g., a
31 query asking for the **PMIX_LOCAL_RANK** of a specified process. Only required when the
32 request is for information on a specific process.

33 **PMIX_NAMESPACE** "pmix.namespace" (**char***)

34 Namespace of the job. Specifies the namespace of the process whose information is being
35 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
36 be accompanied by the **PMIX_RANK** attribute. Only required when the request is for
37 information on a specific process.

1 **PMIX_RANK** "pmix.rank" (pmix_rank_t)
2 Process rank within the job. Specifies the rank of the process whose information is being
3 requested - e.g., a query asking for the **PMIX_LOCAL_RANK** of a specified process. Must
4 be accompanied by the **PMIX_NAMESPACE** attribute. Only required when the request is for
5 information on a specific process.

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.

Optional Attributes

11 The following attributes are optional for host environments that support this operation:

12 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (char*)
13 Request a comma-delimited list of active namespaces.

14 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)
15 Status of a specified, currently executing job.

16 **PMIX_QUERY_QUEUE_LIST** "pmix.qry qlst" (char*)
17 Request a comma-delimited list of scheduler queues.

18 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (TBD)
19 Status of a specified scheduler queue.

20 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)
21 Input namespace of the job whose information is being requested returns (
22 **pmix_data_array_t**) an array of **pmix_proc_info_t**.

23 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lhtable" (char*)
24 Input namespace of the job whose information is being requested returns (
25 **pmix_data_array_t**) an array of **pmix_proc_info_t** for processes in job on same
26 node.

27 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)
28 Return a comma-delimited list of supported spawn attributes.

29 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)
30 Return a comma-delimited list of supported debug attributes.

31 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)
32 Return information on memory usage for the processes indicated in the qualifiers.

33 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
34 Report average values.

35 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)

1 Report minimum and maximum values.

2 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)

3 String identifier of the allocation whose status is being requested.

4 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)

5 Query number of seconds (uint32_t) remaining in allocation for the specified namespace.

6
7 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)

8 URI of the PMIx server to be contacted. Requests the URI of the specified PMIx server's
9 PMIx connection. Defaults to requesting the information for the local PMIx server.

10 **PMIX_PROC_URI** "pmix.puri" (char*)

11 URI containing contact information for a given process. Requests the URI of the specified
12 PMIx server's out-of-band connection. Defaults to requesting the information for the local
13 PMIx server.



14 Description

15 Query information about the system in general. This can include a list of active namespaces,
16 network topology, etc. Also can be used to query node-specific info such as the list of peers
17 executing on a given node. We assume that the host RM will exercise appropriate access control on
18 the information.

19 NOTE: There is no blocking form of this API as the structures passed to query info differ from
20 those for receiving the results.

21 The *status* argument to the callback function indicates if requested data was found or not. An array
22 of `pmix_info_t` will contain each key that was provided and the corresponding value that was
23 found. Requests for keys that are not found will return the key paired with a value of type
24 **PMIX_UNDEF**.

Advice to users

25 The desire to query a list of attributes supported by the implementation and/or the host environment
26 has been expressed and noted. The PMIx community is exploring the possibility and it will likely
27 become available in a future release

Advice to PMIx library implementers

28 Information returned from `PMIx_Query_info_nb` shall be locally cached so that retrieval by
29 subsequent calls to `PMIx_Get` or `PMIx_Query_info_nb` can succeed with minimal overhead.
30 The local cache shall be checked prior to querying the PMIx server and/or the host environment.
31 Queries that include the `PMIX_QUERY_REFRESH_CACHE` attribute shall bypass the local cache
32 and retrieve a new value for the query, refreshing the values in the cache upon return.

1 7.1.3.1 Using `PMIx_Get` vs `PMIx_Query_info_nb`

2 Both `PMIx_Get` and `PMIx_Query_info_nb` can be used to retrieve information about the
3 system. In general, the *get* operation should be used to retrieve:

- 4 • information provided by the host environment at time of job start. This includes information on
5 the number of processes in the job, their location, and possibly their communication endpoints
- 6 • information posted by processes via the `PMIx_Put` function

7 This information is largely considered to be *static*, although this will not necessarily be true for
8 environments supporting dynamic programming models or fault tolerance. Note that the
9 `PMIx_Get` function only accesses information about execution environments - i.e., its scope is
10 limited to values pertaining to a specific `session`, `job`, `application`, process, or node. It
11 cannot be used to obtain information about areas such as the status of queues in the WLM.

12 In contrast, the *query* option should be used to access:

- 13 • system-level information (such as the available WLM queues) that would generally not be
14 included in job-level information provided at job start
- 15 • dynamic information such as application and queue status, and resource utilization statistics.
16 Note that the `PMIX_QUERY_REFRESH_CACHE` attribute must be provided on each query to
17 ensure current data is returned
- 18 • information created post job start, such as process tables
- 19 • information requiring more complex search criteria than supported by the simpler `PMIx_Get`
20 API
- 21 • queries focused on retrieving multi-attribute blocks of data with a single request, thus bypassing
22 the single-key limitation of the `PMIx_Get` API

23 In theory, all information can be accessed via `PMIx_Query_info_nb` as the local cache is
24 typically the same datastore searched by `PMIx_Get`. However, in practice, the overhead
25 associated with the *query* operation may (depending upon implementation) be higher than the
26 simpler *get* operation due to the need to construct and process the more complex `pmix_query_t`
27 structure. Thus, requests for a single key value are likely to be accomplished faster with
28 `PMIx_Get` versus the *query* operation.

29 7.2 Allocation Requests

30 This section defines functionality to request new allocations from the RM, and request
31 modifications to existing allocations. These are primarily used in the following scenarios:

- 32 • *Evolving* applications that dynamically request and return resources as they execute

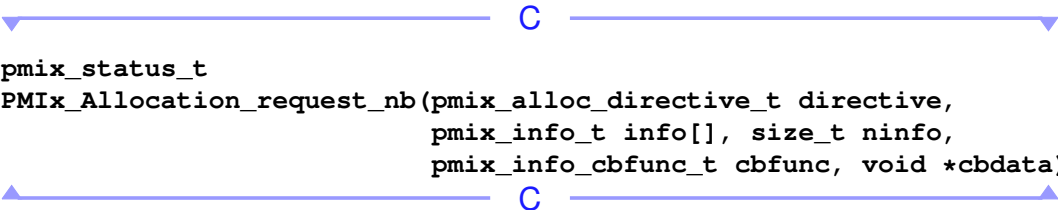
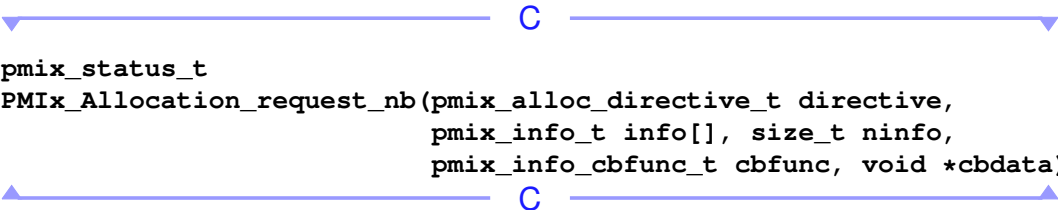
- 1 • *Malleable* environments where the scheduler redirects resources away from executing
- 2 applications for higher priority jobs or load balancing
- 3 • *Resilient* applications that need to request replacement resources in the face of failures
- 4 • *Rigid* jobs where the user has requested a static allocation of resources for a fixed period of time,
- 5 but realizes that they underestimated their required time while executing
- 6 PMIx attempts to address this range of use-cases with a single, flexible API.

7 7.2.1 PMIx_Allocation_request_nb

8 Summary

9 Request an allocation operation from the host resource manager.

10 Format

11 *PMIx v2.0*  C 

```

12 pmix_status_t
13 PMIx_Allocation_request_nb(pmix_alloc_directive_t directive,
14                           pmix_info_t info[], size_t ninfo,
15                           pmix_info_cbfunc_t cbfunc, void *cbdata);

```

- 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)
- 18 **IN cbfunc**
Callback function [pmix_info_cbfunc_t](#) (function reference)
- 19 **IN cbdata**
Data to be passed to the callback function (memory reference)

20 Returns one of the following:

- 21 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
- 22 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback
- 23 function prior to returning from the API.
- 24 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
- 25 returned *success* - the *cbfunc* will *not* be called
- 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
- 28
- 29
- 30
- 31
- 32

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 making the request.

Host environments that implement support for this operation 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. If not given as part of an **pmix_info_t** struct that identifies the involved nodes, then the description will be applied across all nodes in the requestor's allocation.

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

Name of the network.

PMIX_ALLOC_BANDWIDTH "pmix.alloc.bw" (float)

Mbits/sec.

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

2 Quality of service level.



3 **Description**

4 Request an allocation operation from the host resource manager. Several broad categories are
5 envisioned, including the ability to:

- 6 • Request allocation of additional resources, including memory, bandwidth, and compute. This
7 should be accomplished in a non-blocking manner so that the application can continue to
8 progress while waiting for resources to become available. Note that the new allocation will be
9 disjoint from (i.e., not affiliated with) the allocation of the requestor - thus the termination of one
10 allocation will not impact the other.
- 11 • Extend the reservation on currently allocated resources, subject to scheduling availability and
12 priorities. This includes extending the time limit on current resources, and/or requesting
13 additional resources be allocated to the requesting job. Any additional allocated resources will be
14 considered as part of the current allocation, and thus will be released at the same time.
- 15 • Return no-longer-required resources to the scheduler. This includes the “loan” of resources back
16 to the scheduler with a promise to return them upon subsequent request.

17 **7.2.2 PMIx_Job_control_nb**

18 The **PMIx_Job_control_nb** API enables the application and SMS to coordinate the response
19 to failures and other events. This can include requesting termination of the entire job or a subset of
20 processes within a job, but can also be used in combination with other PMIx capabilities (e.g.,
21 allocation support and event notification) for more nuanced responses. For example, an application
22 notified of an incipient over-temperature condition on a node could use the
23 **PMIx_Allocation_request_nb** interface to request replacement nodes while
24 simultaneously using the **PMIx_Job_control_nb** interface to direct that a checkpoint event be
25 delivered to all processes in the application. If replacement resources are not available, the
26 application might use the **PMIx_Job_control_nb** interface to request that the job continue at
27 a lower power setting, perhaps sufficient to avoid the over-temperature failure.

28 The job control API can also be used by an application to register itself as available for preemption
29 when operating in an environment such as a cloud or where incentives, financial or otherwise, are
30 provided to jobs willing to be preempted. Registration can include attributes indicating how many
31 resources are being offered for preemption (e.g., all or only some portion), whether the application
32 will require time to prepare for preemption, etc. Jobs that request a warning will receive an event
33 notifying them of an impending preemption (possibly including information as to the resources that
34 will be taken away, how much time the application will be given prior to being preempted, whether
35 the preemption will be a suspension or full termination, etc.) so they have an opportunity to save
36 their work. Once the application is ready, it calls the provided event completion callback function to
37 indicate that the SMS is free to suspend or terminate it, and can include directives regarding any
38 desired restart.

Summary

Request a job control action.

Format

PMIx v2.0

```
pmix_status_t
PMIx_Job_control_nb(const pmix_proc_t targets[], size_t ntargets,
                   const pmix_info_t directives[], size_t ndirs,
                   pmix_info_cbfunc_t cbfunc, void *cbdata)
```

IN targets

Array of proc structures (array of handles)

IN ntargets

Number of element in the *targets* array (integer)

IN directives

Array of info structures (array of handles)

IN ndirs

Number of element in the *directives* 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 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

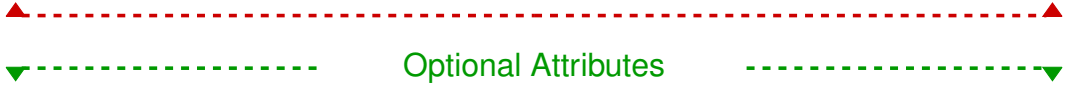
Required Attributes

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 making the request.

Host environments that implement support for this operation are required to support the following attributes:

PMIX_JOB_CTRL_ID "pmix.jctrl.id" (**char***)


- 1 Provide a string identifier for this request.
- 2 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)
- 3 Pause the specified processes.
- 4 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)
- 5 Resume ("un-pause") the specified processes.
- 6 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
- 7 Forcibly terminate the specified processes and cleanup.
- 8 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
- 9 Send given signal to specified processes.
- 10 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
- 11 Politely terminate the specified processes.



Optional Attributes

- 12 The following attributes are optional for host environments that support this operation:
- 13 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
 - 14 Cancel the specified request (**NULL** implies cancel all requests from this requestor).
 - 15 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
 - 16 Restart the specified processes using the given checkpoint ID.
 - 17 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
 - 18 Checkpoint the specified processes and assign the given ID to it.
 - 19 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
 - 20 Use event notification to trigger a process checkpoint.
 - 21 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
 - 22 Use the given signal to trigger a process checkpoint.
 - 23 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
 - 24 Time in seconds to wait for a checkpoint to complete.
 - 25 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
 - 26 "pmix.jctrl.ckmethod" (pmix_data_array_t)
 - 27 Array of **pmix_info_t** declaring each method and value supported by this application.
 - 28 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
 - 29 Regular expression identifying nodes that are to be provisioned.
 - 30 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
 - 31 Name of the image that is to be provisioned.
 - 32 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)

1 Indicate that the job can be pre-empted.



2 **Description**

3 Request a job control action. The *targets* array identifies the processes to which the requested job
4 control action is to be applied. A **NULL** value can be used to indicate all processes in the caller's
5 namespace. The use of **PMIX_RANK_WILDARD** can also be used to indicate that all processes in
6 the given namespace are to be included.

7 The directives are provided as **pmix_info_t** structures in the *directives* array. The callback
8 function provides a *status* to indicate whether or not the request was granted, and to provide some
9 information as to the reason for any denial in the **pmix_info_cbfunc_t** array of
10 **pmix_info_t** structures.

11 **7.3 Process and Job Monitoring**

12 In addition to external faults, a common problem encountered in HPC applications is a failure to
13 make progress due to some internal conflict in the computation. These situations can result in a
14 significant waste of resources as the SMS is unaware of the problem, and thus cannot terminate the
15 job. Various watchdog methods have been developed for detecting this situation, including
16 requiring a periodic “heartbeat” from the application and monitoring a specified file for changes in
17 size and/or modification time.

18 At the request of SMS vendors and members, a monitoring support interface has been included in
19 the PMIx v2 standard. The defined API allows applications to request monitoring, directing what is
20 to be monitored, the frequency of the associated check, whether or not the application is to be
21 notified (via the event notification subsystem) of stall detection, and other characteristics of the
22 operation. In addition, heartbeat and file monitoring methods have been included in the PRI but are
23 active only when requested.

24 **7.3.1 PMIx_Process_monitor_nb**

25 **Summary**

26 Request that application processes be monitored.

1

Format

PMIx v2.0

C

2

`pmix_status_t`

3

`PMIx_Process_monitor_nb(const pmix_info_t *monitor, pmix_status_t error,`

4

`const pmix_info_t directives[], size_t ndirs,`

5

`pmix_info_cbfunc_t cbfunc, void *cbdata)`

C

6

IN monitor

7

info (handle)

8

IN error

9

status (integer)

10

IN directives

11

Array of info structures (array of handles)

12

IN ndirs

13

Number of elements in the *directives* array (integer)

14

IN cbfunc

15

Callback function `pmix_info_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 will be returned in the provided *cbfunc*. Note that the library *must not* invoke the callback function prior to returning from the API.

20

21

22

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

23

24

- 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

25

Optional Attributes

26

The following attributes may be implemented by a PMIx library or by the host environment. If supported by the PMIx server library, then the library must not pass the supported attributes to the host environment. All attributes not directly supported by the server library must be passed to the host environment if it supports this operation, and the library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the requesting process:

27

28

29

30

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

31

32

Provide a string identifier for this request.

33

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

34

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 The *cbfunc* function provides a *status* to indicate whether or not the request was granted, and to
2 provide some information as to the reason for any denial in the `pmix_info_cbfunc_t` array of
3 `pmix_info_t` structures.

4 7.3.2 PMIx_Heartbeat

5 Summary

6 Send a heartbeat to the PMIx server library

7 Format

PMIx v2.0

8 `void PMIx_Heartbeat(void)`

C

C

9 Description

10 A simplified macro wrapping `PMIx_Process_monitor_nb` that sends a heartbeat to the
11 PMIx server library.

12 7.4 Logging

13 The logging interface supports posting information by applications and SMS elements to persistent
14 storage. This function is *not* intended for output of computational results, but rather for reporting
15 status and saving state information such as inserting computation progress reports into the
16 application's SMS job log or error reports to the local syslog.

17 7.4.1 PMIx_Log_nb

18 Summary

19 Log data to a data service.

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

Format

PMIx v2.0

```
pmix_status_t  
PMIx_Log_nb(const pmix_info_t data[], size_t ndata,  
            const pmix_info_t directives[], size_t ndirs,  
            pmix_op_cbfunc_t cbfunc, void *cbdata)
```

- IN data**
Array of info structures (array of handles)
- IN ndata**
Number of elements in the *data* array (**size_t**)
- IN directives**
Array of info 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)

Return codes are one of the following:

- PMIX_SUCCESS** The logging request is valid and is being processed. The resulting status from the operation will be provided in the callback function. 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
- PMIX_ERR_BAD_PARAM** The logging request contains at least one incorrect entry that prevents it from being processed. The callback function will *not* be called.
- PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function. The callback function will *not* be called.

Required Attributes

If the PMIx library does not itself perform this operation, 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 Host environments that implement support for this operation are required to support the following
2 attributes:

3 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)

4 Log string to **stderr**.

5 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)

6 Log string to **stdout**.

7 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)

8 Log data to syslog. Defaults to **ERROR** priority.



▼----- Optional Attributes -----▼

9 The following attributes are optional for host environments that support this operation:

10 **PMIX_LOG_MSG** "pmix.log.msg" (pmix_byte_object_t)

11 Message blob to be sent somewhere.

12 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)

13 Log via email based on **pmix_info_t** containing directives.

14 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)

15 Comma-delimited list of email addresses that are to receive the message.

16 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)

17 Subject line for email.

18 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)

19 Message to be included in email.



20 **Description**

21 Log data subject to the services offered by the host environment. The data to be logged is provided
22 in the *data* array. The (optional) *directives* can be used to direct the choice of logging channel. The
23 callback function will be executed when the log operation has been completed. The *data* and
24 *directives* arrays must be maintained until the callback is provided.

▼----- Advice to users -----▼

25 It is strongly recommended that the **PMIx_Log_nb** API not be used by applications for streaming
26 data as it is not a “performant” transport and can perturb the application since it involves the local
27 PMIx server and host SMS daemon.



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);
```

```

1  IN codes
2      Array of status codes (array of pmix_status_t)
3  IN ncodes
4      Number of elements in the codes array (size_t)
5  IN info
6      Array of info structures (array of handles)
7  IN ninfo
8      Number of elements in the info array (size_t)
9  IN evhdlr
10     Event handler to be called pmix_notification_fn_t (function reference)
11 IN cbfunc
12     Callback function pmix_evhdlr_reg_cbfunc_t (function reference)
13 IN cbdata
14     Data to be passed to the cbfunc callback function (memory reference)

```

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

```

15 PMIX_EVENT_HDLR_NAME "pmix.evname" (char*)
16     String name identifying this handler.
17
18 PMIX_EVENT_HDLR_FIRST "pmix.evfirst" (bool)
19     Invoke this event handler before any other handlers.
20
21 PMIX_EVENT_HDLR_LAST "pmix.evlast" (bool)
22     Invoke this event handler after all other handlers have been called.
23
24 PMIX_EVENT_HDLR_FIRST_IN_CATEGORY "pmix.evfirstcat" (bool)
25     Invoke this event handler before any other handlers in this category.
26
27 PMIX_EVENT_HDLR_LAST_IN_CATEGORY "pmix.evlastcat" (bool)
28     Invoke this event handler after all other handlers in this category have been called.
29
30 PMIX_EVENT_HDLR_BEFORE "pmix.evbefore" (char*)
31     Put this event handler immediately before the one specified in the (char*) value.
32
33 PMIX_EVENT_HDLR_AFTER "pmix.evafter" (char*)
34     Put this event handler immediately after the one specified in the (char*) value.
35
36 PMIX_EVENT_HDLR_PREPEND "pmix.evprepend" (bool)
37     Prepend this handler to the precedence list within its category.
38
39 PMIX_EVENT_HDLR_APPEND "pmix.evappend" (bool)
40     Append this handler to the precedence list within its category.
41
42 PMIX_EVENT_CUSTOM_RANGE "pmix.evrage" (pmix_data_array_t*)

```

1 Array of `pmix_proc_t` defining range of event notification.
2 **PMIX_RANGE** "pmix.range" (`pmix_data_range_t`)
3 Value for calls to publish/lookup/unpublish or for monitoring event notifications.
4 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (`void *`)
5 Object to be returned whenever the registered callback function `cbfunc` is invoked. The
6 object will *only* be returned to the process that registered it.
7 Host environments that implement support for PMIx event notification are required to support the
8 following attributes:
9 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (`pmix_proc_t`)
10 The single process that was affected.
11 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (`pmix_data_array_t*`)
12 Array of `pmix_proc_t` defining affected processes.



▼----- Optional Attributes -----▼

13 Host environments that support PMIx event notification *may* offer notifications for environmental
14 events impacting the job and for SMS events relating to the job. The following attributes are
15 optional for host environments that support this operation:
16 **PMIX_EVENT_TERMINATE_SESSION** "pmix.evterm.sess" (`bool`)
17 The RM intends to terminate this session.
18 **PMIX_EVENT_TERMINATE_JOB** "pmix.evterm.job" (`bool`)
19 The RM intends to terminate this job.
20 **PMIX_EVENT_TERMINATE_NODE** "pmix.evterm.node" (`bool`)
21 The RM intends to terminate all processes on this node.
22 **PMIX_EVENT_TERMINATE_PROC** "pmix.evterm.proc" (`bool`)
23 The RM intends to terminate just this process.
24 **PMIX_EVENT_ACTION_TIMEOUT** "pmix.evtimeout" (`int`)
25 The time in seconds before the RM will execute error response.
26 **PMIX_EVENT_SILENT_TERMINATION** "pmix.evsilentterm" (`bool`)
27 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.

Upon completion, the callback will receive a status based on the following table:

`PMIX_SUCCESS` The event handler was successfully registered - the event handler identifier is returned in the callback.

`PMIX_ERR_BAD_PARAM` One or more of the directives provided in the *info* array was unrecognized.

`PMIX_ERR_NOT_SUPPORTED` The PMIx implementation does not support event notification, or the host SMS does not support notification of the specified event code.

The callback function *must not* be executed prior to returning from the API.

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.

1 **Format**

PMIx v2.0

C

```

2 void
3 PMIx_Deregister_event_handler(size_t evhdlr_ref,
4                               pmix_op_cbfunc_t cbfunc,
5                               void *cbdata);

```

C

- 6 **IN evhdlr_ref**
Event handler ID returned by registration (**size_t**)
- 7
- 8 **IN cbfunc**
Callback function to be executed upon completion of operation **pmix_op_cbfunc_t**
(function reference)
- 9
- 10
- 11 **IN cbdata**
Data to be passed to the cbfunc callback function (memory reference)
- 12

13 **Description**

14 Deregister an event handler. If non-NULL, the provided cbfunc will be called to confirm removal
15 of the designated handler, including a status code as per the following:

- 16 **PMIX_SUCCESS** The event handler was successfully deregistered.
- 17 **PMIX_ERR_BAD_PARAM** The provided *evhdlr_ref* was unrecognized.
- 18 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support event notification.

19 The callback function *must not* be executed prior to returning from the API.

20 **8.1.3 PMIx_Notify_event**

21 **Summary**

22 Report an event for notification via any registered event handler.

23 **Format**

PMIx v2.0

C

```

24 pmix_status_t
25 PMIx_Notify_event(pmix_status_t status,
26                  const pmix_proc_t *source,
27                  pmix_data_range_t range,
28                  pmix_info_t info[], size_t ninfo,
29                  pmix_op_cbfunc_t cbfunc, void *cbdata);

```

1 **IN status**
 2 Status code of the event (`pmix_status_t`)
 3 **IN source**
 4 Pointer to a `pmix_proc_t` identifying the original reporter of the event (handle)
 5 **IN range**
 6 Range across which this notification shall be delivered (`pmix_data_range_t`)
 7 **IN info**
 8 Array of `pmix_info_t` structures containing any further info provided by the originator
 9 of the event (array of handles)
 10 **IN ninfo**
 11 Number of elements in the *info* array (`size_t`)
 12 **IN cbfunc**
 13 Callback function to be executed upon completion of operation `pmix_op_cbfunc_t`
 14 (function reference)
 15 **IN cbdata**
 16 Data to be passed to the cbfunc callback function (memory reference)

17 **PMIX_SUCCESS** The notification request is valid and is being processed. The callback function
 18 will be called when the process-local operation is complete and will provide the resulting
 19 status of that operation. Note that this does *not* reflect the success or failure of delivering the
 20 event to any recipients. The callback function *must not* be executed prior to returning from
 21 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 request contains at least one incorrect entry that prevents it from
 25 being processed. The callback function will *not* be called.

26 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support event notification,
 27 or in the case of a PMIx server calling the API, the range extended beyond the local node and
 28 the host SMS environment does not support event notification. The callback function will
 29 *not* be called.

Required Attributes

30 The following attributes are required to be supported by all PMIx libraries:

31 **PMIX_EVENT_NON_DEFAULT** "pmix.evnondef" (bool)
 32 Event is not to be delivered to default event handlers.

33 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrage" (`pmix_data_array_t*`)
 34 Array of `pmix_proc_t` defining range of event notification.

35 Host environments that implement support for PMIx event notification are required to provide the
 36 following attributes for all events generated by the environment:

37 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (`pmix_proc_t`)

1 The single process that was affected.

2 `PMIX_EVENT_AFFECTED_PROCS` "pmix.evaffected" (`pmix_data_array_t*`)

3 Array of `pmix_proc_t` defining affected processes.

4 Description

5 Report an event for notification via any registered event handler. This function can be called by any
6 PMIx process, including application processes, PMIx servers, and SMS elements. The PMIx server
7 calls this API to report events it detected itself so that the host SMS daemon distribute and handle
8 them, and to pass events given to it by its host down to any attached client processes for processing.
9 Examples might include notification of the failure of another process, detection of an impending
10 node failure due to rising temperatures, or an intent to preempt the application. Events may be
11 locally generated or come from anywhere in the system.

12 Host SMS daemons call the API to pass events down to its embedded PMIx server both for
13 transmittal to local client processes and for the server's own internal processing.

14 Client application processes can call this function to notify the SMS and/or other application
15 processes of an event it encountered. Note that processes are not constrained to report status values
16 defined in the official PMIx standard — any integer value can be used. Thus, applications are free
17 to define their own internal events and use the notification system for their own internal purposes.

▼ Advice to users ▼

18 The callback function will be called upon completion of the `notify_event` function's actions.
19 At that time, any messages required for executing the operation (e.g., to send the notification to the
20 local PMIx server) will have been queued, but may not yet have been transmitted. The caller is
21 required to maintain the input data until the callback function has been executed — the sole purpose
22 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 Support Macros

9 PMIx provides a set of convenience macros for creating, initiating, and releasing data buffers.

10 9.1.1 PMIX_DATA_BUFFER_CREATE

11 Summary

12 Allocate memory for a `pmix_data_buffer_t` object and initialize it

13 Format

14 *PMIx v2.0* ▼  

14 **PMIX_DATA_BUFFER_CREATE**(`buffer`);

▲  

15 **OUT** `buffer`

16 Variable to be assigned the pointer to the allocated `pmix_data_buffer_t` (handle)

17 Description

18 This macro uses *calloc* to allocate memory for the buffer and initialize all fields in it

1 9.1.2 PMIX_DATA_BUFFER_RELEASE

2 Summary

3 Free a `pmix_data_buffer_t` object and the data it contains

4 Format

PMIx v2.0

```
▼ _____ C _____ ▼  
PMIX_DATA_BUFFER_RELEASE(buffer);  
▲ _____ C _____ ▲
```

6 **IN** `buffer`

7 Pointer to the `pmix_data_buffer_t` to be released (handle)

8 Description

9 Free's the data contained in the buffer, and then free's the buffer itself

10 9.1.3 PMIX_DATA_BUFFER_CONSTRUCT

11 Summary

12 Initialize a statically declared `pmix_data_buffer_t` object

13 Format

PMIx v2.0

```
▼ _____ C _____ ▼  
PMIX_DATA_BUFFER_CONSTRUCT(buffer);  
▲ _____ C _____ ▲
```

15 **IN** `buffer`

16 Pointer to the allocated `pmix_data_buffer_t` that is to be initialized (handle)

17 Description

18 Initialize a pre-allocated buffer object

19 9.1.4 PMIX_DATA_BUFFER_DESTRUCT

20 Summary

21 Release the data contained in a `pmix_data_buffer_t` object

1 **Format**

PMIx v2.0

```

▼ _____ C _____ ▼
2 PMIX_DATA_BUFFER_DESTRUCT(buffer);
▲ _____ C _____ ▲

```

3 **IN buffer**

4 Pointer to the `pmix_data_buffer_t` whose data is to be released (handle)

5 **Description**

6 Free's the data contained in a `pmix_data_buffer_t` object

7 **9.1.5 PMIX_DATA_BUFFER_LOAD**

8 **Summary**

9 Load a blob into a `pmix_data_buffer_t` object

10 **Format**

PMIx v2.0

```

▼ _____ C _____ ▼
11 PMIX_DATA_BUFFER_LOAD(buffer, data, size);
▲ _____ C _____ ▲

```

12 **IN buffer**

13 Pointer to a pre-allocated `pmix_data_buffer_t` (handle)

14 **IN data**

15 Pointer to a blob (`char*`)

16 **IN size**

17 Number of bytes in the blob `size_t`

18 **Description**

19 Load the given data into the provided `pmix_data_buffer_t` object, usually done in
20 preparation for unpacking the provided data. Note that the data is *not* copied into the buffer - thus,
21 the blob must not be released until after operations on the buffer have completed.

22 **9.1.6 PMIX_DATA_BUFFER_UNLOAD**

23 **Summary**

24 Unload the data from a `pmix_data_buffer_t` object

1 **Format**

PMIx v2.0

C

2 **PMIX_DATA_BUFFER_UNLOAD**(buffer, data, size);

C

3 **IN** buffer

 Pointer to the `pmix_data_buffer_t` whose data is to be extracted (handle)

4 **OUT** data

 Variable to be assigned the pointer to the extracted blob (`void*`)

5 **OUT** size

 Variable to be assigned the number of bytes in the blob `size_t`

9 **Description**

10 Extract the data in a buffer, assigning the pointer to the data (and the number of bytes in the blob) to
11 the provided variables, usually done to transmit the blob to a remote process for unpacking. The
12 buffer's internal pointer will be set to NULL to protect the data upon buffer destruct or release -
13 thus, the user is responsible for releasing the blob when done with it.

14 **9.2 General Routines**

15 The following routines are provided to support internode transfers in heterogeneous environments.

16 **9.2.1 PMIx_Data_pack**

17 **Summary**

18 Pack one or more values of a specified type into a buffer, usually for transmission to another process

19 **Format**

PMIx v2.0

C

20 **pmix_status_t**

21 **PMIx_Data_pack**(const pmix_proc_t *target,
22 pmix_data_buffer_t *buffer,
23 void *src, int32_t num_vals,
24 pmix_data_type_t type);

- 1 **IN target**
 2 Pointer to a `pmix_proc_t` containing the nspace/rank of the process that will be
 3 unpacking the final buffer. A NULL value may be used to indicate that the target is based on
 4 the same PMIx version as the caller. Note that only the target's nspace is relevant. (handle)
- 5 **IN buffer**
 6 Pointer to a `pmix_data_buffer_t` where the packed data is to be stored (handle)
- 7 **IN src**
 8 Pointer to a location where the data resides. Strings are to be passed as (char **) — i.e., the
 9 caller must pass the address of the pointer to the string as the (void*). This allows the caller
 10 to pass multiple strings in a single call. (memory reference)
- 11 **IN num_vals**
 12 Number of elements pointed to by the `src` pointer. A string value is counted as a single value
 13 regardless of length. The values must be contiguous in memory. Arrays of pointers (e.g.,
 14 string arrays) should be contiguous, although the data pointed to need not be contiguous
 15 across array entries.(`int32_t`)
- 16 **IN type**
 17 The type of the data to be packed (`pmix_data_type_t`)

18 Returns one of the following:

- 19 **PMIX_SUCCESS** The data has been packed as requested
 20 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function.
 21 **PMIX_ERR_BAD_PARAM** The provided buffer or src is **NULL**
 22 **PMIX_ERR_UNKNOWN_DATA_TYPE** The specified data type is not known to this
 23 implementation
 24 **PMIX_ERR_OUT_OF_RESOURCE** Not enough memory to support the operation
 25 **PMIX_ERROR** General error

26 Description

27 The pack function packs one or more values of a specified type into the specified buffer. The buffer
 28 must have already been initialized via the `PMIX_DATA_BUFFER_CREATE` or
 29 `PMIX_DATA_BUFFER_CONSTRUCT` macros — otherwise, `PMIx_Data_pack` will return an
 30 error. Providing an unsupported type flag will likewise be reported as an error.

31 Note that any data to be packed that is not hard type cast (i.e., not type cast to a specific size) may
 32 lose precision when unpacked by a non-homogeneous recipient. The `PMIx_Data_pack` function
 33 will do its best to deal with heterogeneity issues between the packer and unpacker in such cases.
 34 Sending a number larger than can be handled by the recipient will return an error code (generated
 35 upon unpacking) — the error cannot be detected during packing.

36 The namespace of the intended recipient of the packed buffer (i.e., the process that will be
 37 unpacking it) is used solely to resolve any data type differences between PMIx versions. The
 38 recipient must, therefore, be known to the user prior to calling the pack function so that the PMIx

1 library is aware of the version the recipient is using. Note that all processes in a given namespace
2 are *required* to use the same PMIx version — thus, the caller must only know at least one process
3 from the target’s namespace.

4 9.2.2 PMIx_Data_unpack

5 Summary

6 Unpack values from a `pmix_data_buffer_t`

7 Format

PMIx v2.0

C

```
8 pmix_status_t  
9 PMIx_Data_unpack(const pmix_proc_t *source,  
10                 pmix_data_buffer_t *buffer, void *dest,  
11                 int32_t *max_num_values,  
12                 pmix_data_type_t type);  
13
```

C

14 IN source

15 Pointer to a `pmix_proc_t` structure containing the nspace/rank of the process that packed
16 the provided buffer. A NULL value may be used to indicate that the source is based on the
17 same PMIx version as the caller. Note that only the source’s nspace is relevant. (handle)

18 IN buffer

19 A pointer to the buffer from which the value will be extracted. (handle)

20 INOUT dest

21 A pointer to the memory location into which the data is to be stored. Note that these values
22 will be stored contiguously in memory. For strings, this pointer must be to (char**) to
23 provide a means of supporting multiple string operations. The unpack function will allocate
24 memory for each string in the array - the caller must only provide adequate memory for the
25 array of pointers. (void*)

26 INOUT max_num_values

27 The number of values to be unpacked — upon completion, the parameter will be set to the
28 actual number of values unpacked. In most cases, this should match the maximum number
29 provided in the parameters — but in no case will it exceed the value of this parameter. Note
30 that unpacking fewer values than are actually available will leave the buffer in an unpackable
31 state — the function will return an error code to warn of this condition. (int32_t)

32 IN type

33 The type of the data to be unpacked — must be one of the PMIx defined data types (
34 `pmix_data_type_t`)

1 Returns one of the following:

2 **PMIX_SUCCESS** The data has been unpacked as requested

3 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function.

4 **PMIX_ERR_BAD_PARAM** The provided buffer or dest is **NULL**

5 **PMIX_ERR_UNKNOWN_DATA_TYPE** The specified data type is not known to this
6 implementation

7 **PMIX_ERR_OUT_OF_RESOURCE** Not enough memory to support the operation

8 **PMIX_ERROR** General error

9 Description

10 The unpack function unpacks the next value (or values) of a specified type from the given buffer.
11 The buffer must have already been initialized via an **PMIX_DATA_BUFFER_CREATE** or
12 **PMIX_DATA_BUFFER_CONSTRUCT** call (and assumedly filled with some data) — otherwise,
13 the unpack_value function will return an error. Providing an unsupported type flag will likewise be
14 reported as an error, as will specifying a data type that *does not* match the type of the next item in
15 the buffer. An attempt to read beyond the end of the stored data held in the buffer will also return an
16 error.

17 NOTE: it is possible for the buffer to be corrupted and that PMIx will *think* there is a proper
18 variable type at the beginning of an unpack region — but that the value is bogus (e.g., just a byte
19 field in a string array that so happens to have a value that matches the specified data type flag).
20 Therefore, the data type error check is *not* completely safe.

21 Unpacking values is a "nondestructive" process — i.e., the values are not removed from the buffer.
22 It is therefore possible for the caller to re-unpack a value from the same buffer by resetting the
23 unpack_ptr.

24 Warning: The caller is responsible for providing adequate memory storage for the requested data.
25 The user must provide a parameter indicating the maximum number of values that can be unpacked
26 into the allocated memory. If more values exist in the buffer than can fit into the memory storage,
27 then the function will unpack what it can fit into that location and return an error code indicating
28 that the buffer was only partially unpacked.

29 Note that any data that was not hard type cast (i.e., not type cast to a specific size) when packed may
30 lose precision when unpacked by a non-homogeneous recipient. PMIx will do its best to deal with
31 heterogeneity issues between the packer and unpacker in such cases. Sending a number larger than
32 can be handled by the recipient will return an error code generated upon unpacking — these errors
33 cannot be detected during packing.

34 The namespace of the process that packed the buffer is used solely to resolve any data type
35 differences between PMIx versions. The packer must, therefore, be known to the user prior to
36 calling the pack function so that the PMIx library is aware of the version the packer is using. Note
37 that all processes in a given namespace are *required* to use the same PMIx version — thus, the
38 caller must only know at least one process from the packer's namespace.

1 9.2.3 PMIx_Data_copy

2 Summary

3 Copy a data value from one location to another.

4 Format

PMIx v2.0

```
▼ _____ C _____ ▼  
5 pmix_status_t  
6 PMIx_Data_copy(void **dest, void *src,  
7 pmix_data_type_t type);  
▲ _____ C _____ ▲
```

8 **IN dest**

9 The address of a pointer into which the address of the resulting data is to be stored.
10 (void**)

11 **IN src**

12 A pointer to the memory location from which the data is to be copied (handle)

13 **IN type**

14 The type of the data to be copied — must be one of the PMIx defined data types. (
15 [pmix_data_type_t](#))

16 Returns one of the following:

17 [PMIX_SUCCESS](#) The data has been copied as requested

18 [PMIX_ERR_NOT_SUPPORTED](#) The PMIx implementation does not support this function.

19 [PMIX_ERR_BAD_PARAM](#) The provided src or dest is **NULL**

20 [PMIX_ERR_UNKNOWN_DATA_TYPE](#) The specified data type is not known to this
21 implementation

22 [PMIX_ERR_OUT_OF_RESOURCE](#) Not enough memory to support the operation

23 [PMIX_ERROR](#) General error

24 Description

25 Since registered data types can be complex structures, the system needs some way to know how to
26 copy the data from one location to another (e.g., for storage in the registry). This function, which
27 can call other copy functions to build up complex data types, defines the method for making a copy
28 of the specified data type.

29 9.2.4 PMIx_Data_print

30 Summary

31 Pretty-print a data value.

1 **Format**

PMIx v2.0

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.

19 **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 **9.2.5 PMIx_Data_copy_payload**

23 **Summary**

24 Copy a payload from one buffer to another

25 **Format**

PMIx v2.0

C

```
26 pmix_status_t
27 PMIx_Data_copy_payload(pmix_data_buffer_t *dest,
28                       pmix_data_buffer_t *src);
```

1 **IN dest**
2 Pointer to the destination `pmix_data_buffer_t` (handle)
3 **IN src**
4 Pointer to the source `pmix_data_buffer_t` (handle)

5 Returns one of the following:

6 **PMIX_SUCCESS** The data has been copied as requested
7 **PMIX_ERR_BAD_PARAM** The src and dest `pmix_data_buffer_t` types do not match
8 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function.

9 **Description**

10 This function will append a copy of the payload in one buffer into another buffer. Note that this is
11 *not* a destructive procedure — the source buffer’s payload will remain intact, as will any pre-existing
12 payload in the destination’s buffer. Only the unpacked portion of the source payload will be copied.

CHAPTER 10

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.

10.1 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.

10.1.1 PMIx_generate_regex

Summary

16 Generate a regular expression representation of the input string.

Format

PMIx v1.0

```
18 pmix_status_t  
19 PMIx_generate_regex(const char *input, char **regex)
```

IN input

String to process (string)

OUT regex

Regular expression representation of *input* (string)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

1 **Description**

2 Given a comma-separated list of *input* values, generate a regular expression that can be passed
3 down to the PMIx client for parsing. The caller is responsible for free'ing the resulting string.

4 If values have leading zero's, then that is preserved, as are prefix and suffix strings. For example, an
5 input string of

6 `“odin009.org,odin010.org,odin011.org,odin012.org,odin[102-107].org”`
7 will return a regular expression of `“pmix:odin[009-012,102-107].org”`

▼ **Advice to users** ▼

8 The returned regular expression will have a `“pmix:”` at the beginning of the string. This informs
9 the PMIx parser that the string was produced using the PRI's regular expression generator, and thus
10 that same plugin should be used for parsing the string

11 **10.1.2 PMIx_generate_ppn**

12 **Summary**

13 Generate a regular expression representation of the input string.

14 **Format**

PMIx v1.0 ▼ **C** ▼

15 `pmix_status_t PMIx_generate_ppn(const char *input, char **ppn)`

▲ **C** ▲

16 **IN** `input`
17 String to process (string)

18 **OUT** `regex`
19 Regular expression representation of *input* (string)

20 Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Description

The input is expected to consist of a semicolon-separated list of ranges representing the ranks of processes on each node of the job. Thus, an input of "1-4;2-5;8,10,11,12;6,7,9" would generate a regex of "pmix:2x(3);8,10-12;6-7,9"

Advice to users

The returned regular expression will have a "pmix:" at the beginning of the string. This informs the PMIx parser that the string was produced using the PRI's regular expression generator, and thus that same plugin should be used for parsing the string

10.1.3 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 nspace,
                               int nlocalprocs,
                               pmix_info_t info[], size_t ninfo,
                               pmix_op_cbfunc_t cbfunc, void *cbdata)
```

- IN nspace**
namespace (string)
- IN nlocalprocs**
number of local processes (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:

- 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* to be supported by all PMIx libraries:

9 **PMIX_REGISTER_NODATA** "pmix.reg.nodata" (bool)
10 Registration is for this namespace only, do not copy job data - this attribute is not accessed
11 using the **PMIx_Get**

12 Host environments are *required* to provide the following attributes:

- 13 • for the session containing the given namespace:
 - 14 - **PMIX_UNIV_SIZE** "pmix.univ.size" (uint32_t)
15 Number of allocated slots in a session - each slot may or may not be occupied by an
16 executing process. Note that this attribute is the equivalent to the combination of
17 **PMIX_SESSION_INFO_ARRAY** with the **PMIX_NUM_SLOTS** entry in the array - it
18 is included in the Standard for historical reasons.
- 19 • for the given namespace:
 - 20 - **PMIX_JOBID** "pmix.jobid" (char*)
21 Job identifier assigned by the scheduler.
 - 22 - **PMIX_JOB_SIZE** "pmix.job.size" (uint32_t)
23 Total number of processes in this job across all contained applications
 - 24 - **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)
25 Maximum number of processes for this job.
 - 26 - **PMIX_NODE_MAP** "pmix.nmap" (char*)
27 Regular expression of nodes - see 10.1.3.1 for an explanation of its generation.
 - 28 - **PMIX_PROC_MAP** "pmix.pmap" (char*)
29 Regular expression describing processes on each node - see 10.1.3.1 for an explanation
30 of its generation.
- 31 • for its own node:
 - 32 - **PMIX_LOCAL_SIZE** "pmix.local.size" (uint32_t)
33 Number of processes in this job on this node.
 - 34 - **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)

1 Comma-delimited list of ranks on this node within the specified namespace - referenced
2 using [PMIX_RANK_WILDCARD](#) .

- 3 - [PMIX_LOCAL_CPUSSETS](#) "pmix.lcpus" (char*)
4 Colon-delimited cpusets of local peers within the specified namespace - referenced
5 using [PMIX_RANK_WILDCARD](#) .

6 • for each process in the given namespace:

- 7 - [PMIX_RANK](#) "pmix.rank" (pmix_rank_t)
8 Process rank within the job.
- 9 - [PMIX_LOCAL_RANK](#) "pmix.lrank" (uint16_t)
10 Local rank on this node within this job.
- 11 - [PMIX_NODE_RANK](#) "pmix.nrank" (uint16_t)
12 Process rank on this node spanning all jobs.
- 13 - [PMIX_NODEID](#) "pmix.nodeid" (uint32_t)
14 Node identifier where the specified process is located, expressed as the node's index
15 (beginning at zero) in the array resulting from expansion of the [PMIX_NODE_MAP](#)
16 regular expression for the [job](#)

17 If more than one application is included in the namespace, then the host environment is also
18 *required* to provide the following attributes:

19 • for each application:

- 20 - [PMIX_APPNUM](#) "pmix.appnum" (uint32_t)
21 Application number within the job.
- 22 - [PMIX_APPLDR](#) "pmix.aldr" (pmix_rank_t)
23 Lowest rank in this application within this job - referenced using
24 [PMIX_RANK_WILDCARD](#) .
- 25 - [PMIX_APP_SIZE](#) "pmix.app.size" (uint32_t)
26 Number of processes in this application.

27 • for each process:

- 28 - [PMIX_APP_RANK](#) "pmix.apprank" (pmix_rank_t)
29 Process rank within this application.
- 30 - [PMIX_APPNUM](#) "pmix.appnum" (uint32_t)
31 Application number within the job.



Optional Attributes

The following attributes *may* be provided by host environments:

- for the session containing the given namespace:
 - **PMIX_SESSION_ID** "pmix.session.id" (uint32_t)
Session identifier - referenced using **PMIX_RANK_WILDCARD** .
- for the given namespace:
 - **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (char*)
Name of the namespace to use for this PMIx server.
 - **PMIX_SERVER_RANK** "pmix.srv.rank" (pmix_rank_t)
Rank of this PMIx server
 - **PMIX_NPROC_OFFSET** "pmix.offset" (pmix_rank_t)
Starting global rank of this job - referenced using **PMIX_RANK_WILDCARD** .
 - **PMIX_ALLOCATED_NODELIST** "pmix.alist" (char*)
Comma-delimited list of all nodes in this allocation regardless of whether or not they currently host processes - referenced using **PMIX_RANK_WILDCARD** .
 - **PMIX_JOB_NUM_APPS** "pmix.job.napps" (uint32_t)
Number of applications in this job.
 - **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
- for its own node:
 - **PMIX_AVAIL_PHYS_MEMORY** "pmix.pmem" (uint64_t)
Total available physical memory on this node.
 - **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*)

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

- 2 - **PMIX_LOCALLDR** "pmix.lldr" (pmix_rank_t)
3 Lowest rank on this node within this job - referenced using **PMIX_RANK_WILDCARD** .
4
- 5 - **PMIX_NODE_SIZE** "pmix.node.size" (uint32_t)
6 Number of processes across all jobs on this node.
- 7 - **PMIX_LOCAL_PROCS** "pmix.lprocs" (pmix_proc_t array)
8 Array of **pmix_proc_t** of all processes on the specified node - referenced using
9 **PMIX_RANK_WILDCARD** .
- 10 • for each process in the given namespace:
 - 11 - **PMIX_PROCID** "pmix.procid" (pmix_proc_t)
12 Process identifier
 - 13 - **PMIX_GLOBAL_RANK** "pmix.grank" (pmix_rank_t)
14 Process rank spanning across all jobs in this session.
 - 15 - **PMIX_HOSTNAME** "pmix.hname" (char*)
16 Name of the host where the specified process is running.

17 Attributes not directly provided by the host environment *may* be derived by the PMIx server library
18 from other required information and included in the data made available to the server library's
19 clients.



20 Description

21 Pass job-related information to the PMIx server library for distribution to local client processes.



Advice to PMIx server hosts

22 Host environments are *required* to execute this operation prior to starting any local application
23 process within the given namespace.

24 The PMIx server must register *all* namespaces that will participate in collective operations with
25 local processes. This means that the server must register a namespace even if it will not host any
26 local processes from within that namespace *if* any local process of another namespace might at
27 some point perform an operation involving one or more processes from the new namespace. This is
28 necessary so that the collective operation can identify the participants and know when it is locally
29 complete.

30 The caller must also provide the number of local processes that will be launched within this
31 namespace. This is required for the PMIx server library to correctly handle collectives as a
32 collective operation call can occur before all the local processes have been started.



Advice to users

1 The number of local processes for any given namespace is generally fixed at the time of application
2 launch. Calls to `PMIx_Spawn` result in processes launched in their own namespace, not that of
3 their parent. However, it is possible for processes to *migrate* to another node via a call to
4 `PMIx_Job_control_nb`, thus resulting in a change to the number of local processes on both
5 the initial node and the node to which the process moved. It is therefore *critical* that applications
6 not migrate processes without first ensuring that PMIx-based collective operations are not in
7 progress, and that no such operations be initiated until process migration has completed.

8 10.1.3.1 Assembling the registration information

9 The following description is not intended to represent the actual layout of information in a given
10 PMIx library. Instead, it describes how information provided in the *info* parameter of the
11 `PMIx_server_register_namespace` shall be organized for proper processing by a PMIx server
12 library. The ordering of the various information elements is arbitrary - they are presented in a
13 top-down hierarchical form solely for clarity in reading.

Advice to PMIx server hosts

14 Creating the *info* array of data requires knowing in advance the number of elements required for the
15 array. This can be difficult to compute and somewhat fragile in practice. One method for resolving
16 the problem is to create a linked list of objects, each containing a single `pmix_info_t` structure.
17 Allocation and manipulation of the list can then be accomplished using existing standard methods.
18 Upon completion, the final *info* array can be allocated based on the number of elements on the list,
19 and then the values in the list object `pmix_info_t` structures transferred to the corresponding
20 array element utilizing the `PMIX_INFO_XFER` macro.

21 A common building block used in several areas is the construction of a regular expression
22 identifying the nodes involved in that area - e.g., the nodes in a `session` or `job`. PMIx provides
23 several tools to facilitate this operation, beginning by constructing an argv-like array of node
24 names. This array is then passed to the `PMIx_generate_regex` function to create a regular
25 expression parseable by the PMIx server library, as shown below:

```

1  char **nodes = NULL;
2  char *nodelist;
3  char *regex;
4  size_t n;
5  pmix_status_t rc;
6  pmix_info_t info;
7
8  /* loop over an array of nodes, adding each
9   * name to the array */
10 for (n=0; n < num_nodes; n++)
11     /* filter the nodes to ignore those not included
12      * in the target range (session, job, etc.). In
13      * this example, all nodes are accepted */
14     PMIX_ARGV_APPEND(&nodes, node[n]->name);
15
16
17 /* join into a comma-delimited string */
18 nodelist = PMIX_ARGV_JOIN(nodes, ',');
19
20 /* release the array */
21 PMIX_ARGV_FREE(nodes);
22
23 /* generate regex */
24 rc = PMIX_generate_regex(nodelist, &regex);
25
26 /* release list */
27 free(nodelist);
28
29 /* pass the regex as the value to the PMIX_NODE_MAP key */
30 PMIX_INFO_LOAD(&info, PMIX_NODE_MAP, regex, PMIX_STRING);
31 /* release the regex */
32 free(regex);
33

```

34 Changing the filter criteria allows the construction of node maps for any level of information.

35 A similar method is used to construct the map of processes on each node from the namespace being
36 registered. This may be done for each information level of interest (e.g., to identify the process map
37 for the entire **job** or for each **application** in the job) by changing the search criteria. An
38 example is shown below for the case of creating the process map for a **job** :

```

1  char **ndppn;
2  char rank[30];
3  char **ppnarray = NULL;
4  char *ppn;
5  char *localranks;
6  char *regex;
7  size_t n, m;
8  pmix_status_t rc;
9  pmix_info_t info;
10
11 /* loop over an array of nodes */
12 for (n=0; n < num_nodes; n++)
13     /* for each node, construct an array of ranks on that node */
14     ndppn = NULL;
15     for (m=0; m < node[n]->num_procs; m++)
16         /* ignore processes that are not part of the target job */
17         if (!PMIX_CHECK_NAMESPACE(targetjob,node[n]->proc[m].nspace))
18             continue;
19
20         snprintf(rank, 30, "%d", node[n]->proc[m].rank);
21         PMIX_ARGV_APPEND(&ndppn, rank);
22
23     /* convert the array into a comma-delimited string of ranks */
24     localranks = PMIX_ARGV_JOIN(ndppn, ',');
25     /* release the local array */
26     PMIX_ARGV_FREE(ndppn);
27     /* add this node's contribution to the overall array */
28     PMIX_ARGV_APPEND(&ppnarray, localranks);
29     /* release the local list */
30     free(localranks);
31
32
33 /* join into a semicolon-delimited string */
34 ppn = PMIX_ARGV_JOIN(ppnarray, ';');
35
36 /* release the array */
37 PMIX_ARGV_FREE(ppnarray);
38
39 /* generate ppn regex */
40 rc = PMIx_generate_ppn(ppn, &regex);
41
42 /* release list */

```

```

1  free(ppn);
2
3  /* pass the regex as the value to the PMIX_PROC_MAP key */
4  PMIX_INFO_LOAD(&info, PMIX_PROC_MAP, regex, PMIX_STRING);
5  /* release the regex */
6  free(regex);
7

```



8 Note that the **PMIX_NODE_MAP** and **PMIX_PROC_MAP** attributes are linked in that the order of
9 entries in the process map must match the ordering of nodes in the node map - i.e., there is no
10 provision in the PMIx process map regular expression generator/parser pair supporting an
11 out-of-order node or a node that has no corresponding process map entry (e.g., a node with no
12 processes on it). Armed with these tools, the registration *info* array can be constructed as follows:

- 13 • Session-level information includes all session-specific values. In many cases, only two values (
14 **PMIX_SESSION_ID** and **PMIX_UNIV_SIZE**) are included in the registration array. Since
15 both of these values are session-specific, they can be specified independently - i.e., in their own
16 **pmix_info_t** elements of the *info* array. Alternatively, they can be provided as a
17 **pmix_data_array_t** array of **pmix_info_t** using the **PMIX_SESSION_INFO_ARRAY**
18 attribute and identified by including the **PMIX_SESSION_ID** attribute in the array - this is
19 *required* in cases where non-specific attributes (e.g., **PMIX_NUM_NODES** or
20 **PMIX_NODE_MAP**) are passed to describe aspects of the session. Note that the node map can
21 include nodes not used by the job being registered as no corresponding process map is specified.

22 The *info* array at this point might look like (where the labels identify the corresponding attribute
23 - e.g., “Session ID” corresponds to the **PMIX_SESSION_ID** attribute):

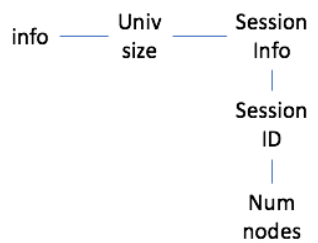


Figure 10.1.: Session-level information elements

- 24 • Job-level information includes all job-specific values such as **PMIX_JOB_SIZE** ,
25 **PMIX_JOB_NUM_APPS** , and **PMIX_JOBID** . Since each invocation of
26 **PMIx_server_register_nspace** describes a single **job** , job-specific values can be
27 specified independently - i.e., in their own **pmix_info_t** elements of the *info* array.
28 Alternatively, they can be provided as a **pmix_data_array_t** array of **pmix_info_t**
29 identified by the **PMIX_JOB_INFO_ARRAY** attribute - this is *required* in cases where

1 non-specific attributes (e.g., `PMIX_NODE_MAP`) are passed to describe aspects of the job. Note
 2 that since the invocation only involves a single namespace, there is no need to include the
 3 `PMIX_NAMESPACE` attribute in the array.

4 Upon conclusion of this step, the *info* array might look like:

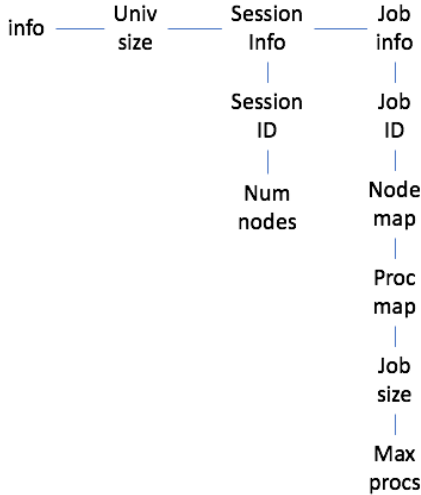


Figure 10.2.: Job-level information elements

5 Note that in this example, `PMIX_NUM_NODES` is not required as that information is contained
 6 in the `PMIX_NODE_MAP` attribute. Similarly, `PMIX_JOB_SIZE` is not technically required as
 7 that information is contained in the `PMIX_PROC_MAP` when combined with the corresponding
 8 node map - however, there is no issue with including the job size as a separate entry.

- 9 • Application-level information includes all application-specific values such as `PMIX_APP_SIZE`
 10 and `PMIX_APPLDR` . If the `job` contains only a single `application` , then the
 11 application-specific values can be specified independently - i.e., in their own `pmix_info_t`
 12 elements of the *info* array - or as a `pmix_data_array_t` array of `pmix_info_t` using the
 13 `PMIX_APP_INFO_ARRAY` attribute and identified by including the `PMIX_APPNUM` attribute
 14 in the array. Use of the array format is *required* in cases where non-specific attributes (e.g.,
 15 `PMIX_NODE_MAP`) are passed to describe aspects of the application.

16 However, in the case of a job consisting of multiple applications, all application-specific values
 17 for each application *must* be provided using the `PMIX_APP_INFO_ARRAY` format, each
 18 identified by its `PMIX_APPNUM` value.

19 Upon conclusion of this step, the *info* array might look like that shown in 10.3, assuming there
 20 are two applications in the job being registered:

- 21 • Process-level information includes an entry for each process in the job being registered, each
 22 entry marked with the `PMIX_PROC_DATA` attribute. The `rank` of the process *must* be the first

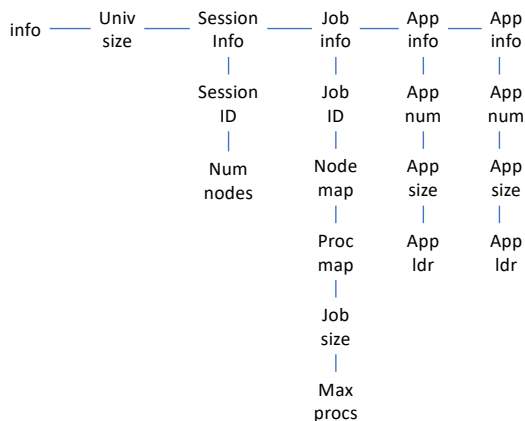


Figure 10.3.: Application-level information elements

1 entry in the array - this provides efficiency when storing the data. Upon conclusion of this step,
 2 the *info* array might look like the diagram in 10.4:

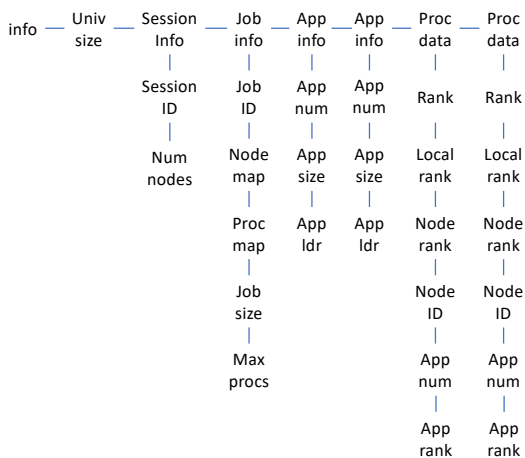


Figure 10.4.: Process-level information elements

- 3
- Node-level information only includes values describing the local node - i.e., it does not include
 4 information about other nodes in the job or session. In many cases, the values included in this
 5 level are unique to it and can be specified independently - i.e., in their own `pmix_info_t`
 6 elements of the *info* array. Alternatively, they can be provided as a `pmix_data_array_t`
 7 array of `pmix_info_t` using the `PMIX_NODE_INFO_ARRAY` attribute - this is *required* in
 8 cases where non-specific attributes are passed to describe aspects of the node.

1 The node-level information requires two elements that must be constructed in a manner similar to
2 that used for the node map. The `PMIX_LOCAL_PEERS` value is computed based on the
3 processes on the local node, filtered to select those from the job being registered, as shown below
4 using the tools provided by PMIx:

C

```
5 char **ndppn = NULL;
6 char rank[30];
7 char *localranks;
8 size_t m;
9 pmix_info_t info;
10
11 for (m=0; m < mynode->num_procs; m++)
12     /* ignore processes that are not part of the target job */
13     if (!PMIX_CHECK_NAMESPACE(targetjob, mynode->proc[m].nspace))
14         continue;
15
16     snprintf(rank, 30, "%d", mynode->proc[m].rank);
17     PMIX_ARGV_APPEND(&ndppn, rank);
18
19     /* convert the array into a comma-delimited string of ranks */
20     localranks = PMIX_ARGV_JOIN(ndppn, ',');
21     /* release the local array */
22     PMIX_ARGV_FREE(ndppn);
23
24     /* pass the string as the value to the PMIX_LOCAL_PEERS key */
25     PMIX_INFO_LOAD(&info, PMIX_LOCAL_PEERS, localranks, PMIX_STRING);
26     /* release the list */
27     free(localranks);
28
```

C

29 The `PMIX_LOCAL_CPUSSETS` value is constructed in a similar manner. In the provided
30 example, it is assumed that the Hardware Locality (HWLOC) cpuset representation (a
31 comma-delimited string of processor IDs) of the processors assigned to each process has
32 previously been generated and stored on the process description. Thus, the value can be
33 constructed as shown below:

C

```
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
```

C

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 [10.5](#):

25 10.1.4 PMIx_server_deregister_namespace

26 Summary

27 Deregister a namespace.

28 Format

PMIx v1.0

C

```
29 void PMIx_server_deregister_namespace(const pmix_namespace_t nspace,
30                                       pmix_op_cbfunc_t cbfunc, void *cbdata)
```

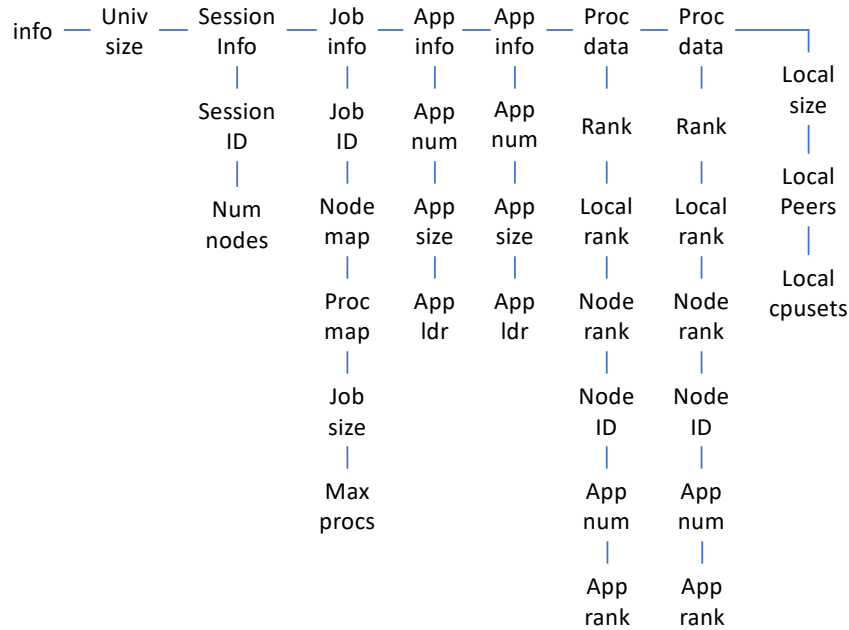


Figure 10.5.: Final information array

C

- 1 **IN** `nspc`
- 2 Namespace (string)
- 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 **Description**

8 Deregister the specified `nspc` and purge all objects relating to it, including any client information
 9 from that namespace. This is intended to support persistent PMIx servers by providing an
 10 opportunity for the host RM to tell the PMIx server library to release all memory for a completed
 11 job. Note that the library *must not* invoke the callback function prior to returning from the API.

12 **10.1.5 PMIx_server_register_client**

13 **Summary**

14 Register a client process with the PMIx server library.

1

Format

PMIx v1.0

C

2

```

pmix_status_t
PMIx_server_register_client(const pmix_proc_t *proc,
                           uid_t uid, gid_t gid,
                           void *server_object,
                           pmix_op_cbfunc_t cbfunc, void *cbdata)

```

5

6

C

7

IN `proc`

8

`pmix_proc_t` structure (handle)

9

IN `uid`

user id (integer)

10

11

IN `gid`

group id (integer)

12

13

IN `server_object`

(memory reference)

14

15

IN `cbfunc`

Callback function `pmix_op_cbfunc_t` (function reference)

16

17

IN `cbdata`

Data to be passed to the callback function (memory reference)

18

19

Returns one of the following:

20

- **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.

21

22

23

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the `cbfunc` will *not* be called

24

25

- 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

26

27

Description

28

Register a client process with the PMIx server library.

29

The host server can also, if it desires, provide an object it wishes to be returned when a server function is called that relates to a specific process. For example, the host server may have an object that tracks the specific client. Passing the object to the library allows the library to provide that object to the host server during subsequent calls related to that client, such as a `pmix_server_client_connected_fn_t` function. This allows the host server to access the object without performing a lookup based on the client's namespace and rank.

30

31

32

33

34

Advice to PMIx server hosts

1 Host environments are *required* to execute this operation prior to starting the client process. The
2 expected user ID and group ID of the child process allows the server library to properly authenticate
3 clients as they connect by requiring the two values to match. Accordingly, the detected user and
4 group ID's of the connecting process are not included in the
5 `pmix_server_client_connected_fn_t` server module function.

Advice to PMIx library implementers

6 For security purposes, the PMIx server library should check the user and group ID's of a
7 connecting process against those provided for the declared client process identifier via the
8 `PMIx_server_register_client` prior to completing the connection.

10.1.6 `PMIx_server_deregister_client`

Summary

Deregister a client and purge all data relating to it.

Format

PMIx v1.0

```
void  
PMIx_server_deregister_client(const pmix_proc_t *proc,  
                             pmix_op_cbfunc_t cbfunc, void *cbdata)
```

IN `proc`
 `pmix_proc_t` structure (handle)
IN `cbfunc`
 Callback function `pmix_op_cbfunc_t` (function reference)
IN `cbdata`
 Data to be passed to the callback function (memory reference)

Description

The `PMIx_server_deregister_namespace` API will delete all client information for that namespace. The PMIx server library will automatically perform that operation upon disconnect of all local clients. This API is therefore intended primarily for use in exception cases, but can be called in non-exception cases if desired. Note that the library *must not* invoke the callback function prior to returning from the API.

1

Format

PMIx v1.0

2

```

pmix_status_t PMIx_server_dmodex_request(const pmix_proc_t *proc,
                                         pmix_dmodex_response_fn_t cbfunc,
                                         void *cbdata)

```

3

4

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 will be returned in the provided `cbfunc`. Note that the library *must not* invoke the callback function prior to returning from the API.

13

14

- a PMIx error constant indicating an error in the input - the `cbfunc` will *not* be called

16

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

While direct modex allows for faster launch times by eliminating the barrier operation, per-peer retrieval of posted information is less efficient. Optimizations can be implemented - e.g., by returning posted information from all processes on a node upon first request - but in general direct modex remains best suited for sparsely connected applications.

10.1.9 PMIx_server_setup_application

Summary

Provide a function by which the resource manager can request application-specific setup data prior to launch of an application.

Format

PMIx v2.0

C

```
pmix_status_t
PMIx_server_setup_application(const pmix_namespace_t nspace,
                             pmix_info_t info[], size_t ninfo,
                             pmix_setup_application_cbfunc_t cbfunc,
                             void *cbdata)
```

C

IN nspace
namespace (string)

IN info
Array of info structures (array of handles)

IN ninfo
Number of elements in the *info* array (integer)

IN cbfunc
Callback function [pmix_setup_application_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 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 either an error in the input - the *cbfunc* will *not* be called

1 **Description**

2 Provide a function by which the RM can request application-specific setup data (e.g., environmental
3 variables, fabric configuration and security credentials) from supporting PMIx server library
4 subsystems prior to initiating launch of an application.

▼————— **Advice to PMIx server hosts** —————▼

5 Host environments are *required* to execute this operation prior to launching an application.



6 This is defined as a non-blocking operation in case contributing subsystems need to perform some
7 potentially time consuming action (e.g., query a remote service) before responding. The returned
8 data must be distributed by the RM and subsequently delivered to the local PMIx server on each
9 node where application processes will execute prior to initiating execution of those processes.

10 In the callback function, the returned *info* array is owned by the PMIx server library and will be
11 free'd when the provided *cbfunc* is called.

▼————— **Advice to PMIx library implementers** —————▼

12 Support for harvesting of environmental variables and providing of local configuration information
13 by the PMIx implementation is optional.



14 **10.1.10 PMIx_server_setup_local_support**

15 **Summary**

16 Provide a function by which the local PMIx server can perform any application-specific operations
17 prior to spawning local clients of a given application.

1

Format

PMIx v2.0

C

2

pmix_status_t

3

```

PMIx_server_setup_local_support(const pmix_namespace_t nspace,
                                pmix_info_t info[], size_t ninfo,
                                pmix_op_cbfunc_t cbfunc,
                                void *cbdata);

```

6

C

7

IN nspace

8

Namespace (string)

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 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 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.

19

20

21

- [PMIX_OPERATION_SUCCEEDED](#) , indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

22

23

- 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

24

25

Description

26

Provide a function by which the local PMIx server can perform any application-specific operations prior to spawning local clients of a given application. For example, a network library might need to setup the local driver for “instant on” addressing. The data provided in the *info* array is the data provided to there host RM from the a call to [PMIx_server_setup_application](#) .

27

28

29

Advice to PMIx server hosts

30

Host environments are *required* to execute this operation prior to starting any local application processes from the specified namespace.

31

1 10.2 Server Function Pointers

2 PMIx utilizes a "function-shipping" approach to support for implementing the server-side of the
3 protocol. This method allows RMs to implement the server without being burdened with PMIx
4 internal details. When a request is received from the client, the corresponding server function will
5 be called with the information.

6 Any functions not supported by the RM can be indicated by a **NULL** for the function pointer. Client
7 calls to such functions will return a **PMIX_ERR_NOT_SUPPORTED** status.

8 The host RM will provide the function pointers in a **pmix_server_module_t** structure passed
9 to **PMIx_server_init** . That module structure and associated function references are defined
10 in this section.

▼ Advice to PMIx server hosts ▼

11 For performance purposes, the host server is required to return as quickly as possible from all
12 functions. Execution of the function is thus to be done asynchronously so as to allow the PMIx
13 server support library to handle multiple client requests as quickly and scalably as possible.

14 *All* data passed to the host server functions is "owned" by the PMIX server support library and
15 *MUST NOT* be free'd. Data returned by the host server via callback function is owned by the host
16 server, which is free to release it upon return from the callback



17 10.2.1 pmix_server_module_t Module

18 Summary

19 List of function pointers that a PMIx server passes to **PMIx_server_init** during startup.

20 Format

```

1 typedef struct pmix_server_module_2_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    pmix_server_module_t;

```

26 10.2.2 pmix_server_client_connected_fn_t

27 Summary

28 Notify the host server that a client connected to this server.

29 Format

PMIx v1.0

```

30 typedef pmix_status_t (*pmix_server_client_connected_fn_t) (
31     const pmix_proc_t *proc,
32     void* server_object,
33     pmix_op_cbfunc_t cbfunc,
34     void *cbdata)

```

1 **IN** `proc`
 2 `pmix_proc_t` structure (handle)
 3 **IN** `server_object`
 4 object reference (memory reference)
 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 Description

18 Notify the host environment that a client has called **PMIx_Init**. Note that the client will be in a
 19 blocked state until the host server executes the callback function, thus allowing the PMIx server
 20 support library to release the client. The `server_object` parameter will be the value of the
 21 `server_object` parameter passed to **PMIx_server_register_client** by the host server
 22 when registering the connecting client. If provided, an implementation of
 23 `pmix_server_client_connected_fn_t` is only required to call the callback function
 24 designated. A host server can choose to not be notified when clients connect by setting
 25 `pmix_server_client_connected_fn_t` to **NULL**.

26 It is possible that only a subset of the clients in a namespace call **PMIx_Init**. The server's
 27 `pmix_server_client_connected_fn_t` implementation should not depend on being
 28 called once per rank in a namespace or delay calling the callback function until all ranks have
 29 connected. However, if a rank makes any PMIx calls, it must first call **PMIx_Init** and therefore
 30 the server's `pmix_server_client_connected_fn_t` will be called before any other
 31 server functions specific to the rank.

Advice to PMIx server hosts

32 This operation is an opportunity for a host environment to update the status of the ranks it manages.
 33 It is also a convenient and well defined time to perform initialization necessary to support further
 34 calls into the server related to that rank.

1 10.2.3 pmix_server_client_finalized_fn_t

2 Summary

3 Notify the host environment that a client called `PMIx_Finalize`.

4 Format

PMIx v1.0

```
▼ C ▼
5 typedef pmix_status_t (*pmix_server_client_finalized_fn_t) (
6     const pmix_proc_t *proc,
7     void* server_object,
8     pmix_op_cbfunc_t cbfunc,
9     void *cbdata)
▲ C ▲
```

10 **IN** `proc`
11 `pmix_proc_t` structure (handle)
12 **IN** `server_object`
13 object reference (memory reference)
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 Notify the host environment that a client called `PMIx_Finalize`. Note that the client will be in
28 a blocked state until the host server executes the callback function, thus allowing the PMIx server
29 support library to release the client. The `server_object` parameter will be the value of the
30 `server_object` parameter passed to `PMIx_server_register_client` by the host server
31 when registering the connecting client. If provided, an implementation of
32 `pmix_server_client_finalized_fn_t` is only required to call the callback function

1 designated. A host server can choose to not be notified when clients finalize by setting
2 `pmix_server_client_finalized_fn_t` to `NULL`.

3 Note that the host server is only being informed that the client has called `PMIx_Finalize`. The
4 client might not have exited. If a client exits without calling `PMIx_Finalize`, the server support
5 library will not call the `pmix_server_client_finalized_fn_t` implementation.

Advice to PMIx server hosts

6 This operation is an opportunity for a host server to update the status of the tasks it manages. It is
7 also a convenient and well defined time to release resources used to support that client.

8 10.2.4 `pmix_server_abort_fn_t`

9 Summary

10 Notify the host environment that a local client called `PMIx_Abort`.

11 Format

PMIx v1.0

C

```
12 typedef pmix_status_t (*pmix_server_abort_fn_t) (  
13     const pmix_proc_t *proc,  
14     void *server_object,  
15     int status,  
16     const char msg[],  
17     pmix_proc_t procs[],  
18     size_t nprocs,  
19     pmix_op_cbfunc_t cbfunc,  
20     void *cbdata)
```

C

21 **IN** `proc`
22 `pmix_proc_t` structure identifying the process requesting the abort (handle)
23 **IN** `server_object`
24 object reference (memory reference)
25 **IN** `status`
26 exit status (integer)
27 **IN** `msg`
28 exit status message (string)
29 **IN** `procs`
30 Array of `pmix_proc_t` structures identifying the processes to be terminated (array of
31 handles)

1 **IN** **nprocs**
2 Number of elements in the *procs* 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 host *must not* invoke the callback function
10 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

15 **Description**

16 A local client called **PMIx_Abort** . Note that the client will be in a blocked state until the host
17 server executes the callback function, thus allowing the PMIx server library to release the client.
18 The array of *procs* indicates which processes are to be terminated. A **NULL** indicates that all
19 processes in the client's namespace are to be terminated.

20 **10.2.5 pmix_server_fence_nb_fn_t**

21 **Summary**

22 At least one client called either **PMIx_Fence** or **PMIx_Fence_nb** .

23 **Format**

PMIx v1.0

C

```
24       typedef pmix_status_t (*pmix_server_fence_nb_fn_t) (  
25                            const pmix_proc_t procs[],  
26                            size_t nprocs,  
27                            const pmix_info_t info[],  
28                            size_t ninfo,  
29                            char *data, size_t ndata,  
30                            pmix_modex_cbfunc_t cbfunc,  
31                            void *cbdata)
```

1 **IN** **procs**
 2 Array of **pmix_proc_t** structures identifying operation participants(array of handles)
 3 **IN** **nprocs**
 4 Number of elements in the *procs* array (integer)
 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** **data**
 10 (string)
 11 **IN** **ndata**
 12 (integer)
 13 **IN** **cbfunc**
 14 Callback function **pmix_modex_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 by the host environment - result
 19 will be returned in the provided *cbfunc*. Note that the host *must not* invoke the callback function
 20 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

Required Attributes

25 PMIx libraries are required to pass any provided attributes to the host environment for processing.

26 The following attributes are required to be supported by all host environments:

27 **PMIX_COLLECT_DATA** "pmix.collect" (bool)
 28 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_fencenb_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.

1 The provided data is to be collectively shared with all PMIx servers involved in the fence operation,
2 and returned in the modex *cbfunc*. A **NULL** data value indicates that the local processes had no data
3 to contribute.

4 The array of *info* structs is used to pass user-requested options to the server. This can include
5 directives as to the algorithm to be used to execute the fence operation. The directives are optional
6 *unless* the **PMIX_INFO_REQD** flag has been set - in such cases, the host RM is required to return
7 an error if the directive cannot be met.

8 10.2.6 pmix_server_dmodex_req_fn_t

9 Summary

10 Used by the PMIx server to request its local host contact the PMIx server on the remote node that
11 hosts the specified proc to obtain and return a direct modex blob for that proc.

12 Format

PMIx v1.0

```
13 typedef pmix_status_t (*pmix_server_dmodex_req_fn_t) (  
14     const pmix_proc_t *proc,  
15     const pmix_info_t info[],  
16     size_t ninfo,  
17     pmix_modex_cbfunc_t cbfunc,  
18     void *cbdata)
```

19 IN **proc**

20 `pmix_proc_t` structure identifying the process whose data is being requested (handle)

21 IN **info**

22 Array of info structures (array of handles)

23 IN **ninfo**

24 Number of elements in the *info* array (integer)

25 IN **cbfunc**

26 Callback function `pmix_modex_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 host *must not* invoke the callback function
32 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.

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.

10.2.7 pmix_server_publish_fn_t

Summary

Publish data per the PMIx API specification.

1

Format

PMIx v1.0

C

2

```

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)

```

3

4

5

6

7

C

8

IN `proc`

9

`pmix_proc_t` structure of the process publishing the data (handle)

10

IN `info`

11

Array of info structures (array of handles)

12

IN `ninfo`

13

Number of elements 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 will be returned in the provided `cbfunc`. Note that the host *must not* invoke the callback function prior to returning from the API.

20

21

22

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the `cbfunc` will *not* be called

23

24

- 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

25

----- Required Attributes -----

26

PMIx libraries are required to pass any provided attributes to the host environment for processing.

27

In addition, the following attributes are required to be included in the passed `info` array:

28

PMIX_USERID "`pmix.euid`" (`uint32_t`)

29

Effective user id.

30

PMIX_GRPID "`pmix.egid`" (`uint32_t`)

31

Effective group id.

32

Host environments that implement this entry point are required to support the following attributes:

33

PMIX_RANGE "`pmix.range`" (`pmix_data_range_t`)

34

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

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

2 Value for calls to **PMIx_Publish** .

3 **Optional Attributes**

3 The following attributes are optional for host environments that support this operation:

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

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

8 **Description**

9 Publish data per the **PMIx_Publish** specification. The callback is to be executed upon
10 completion of the operation. The default data range is left to the host environment, but expected to
11 be **PMIX_SESSION** , and the default persistence **PMIX_PERSIST_SESSION** or their
12 equivalent. These values can be specified by including the respective attributed in the *info* array.

13 The persistence indicates how long the server should retain the data.

14 **Advice to PMIx server hosts**

15 The host environment is not required to guarantee support for any specific range - i.e., the
16 environment does not need to return an error if the data store doesn't support a specified range so
17 long as it is covered by some internally defined range. However, the server must return an error (a)
18 if the key is duplicative within the storage range, and (b) if the server does not allow overwriting of
19 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.

20 The **PMIX_USERID** and **PMIX_GRPID** of the publishing process will be provided to support
21 authorization-based access to published information and must be returned on any subsequent
22 lookup request.

23 **10.2.8 pmix_server_lookup_fn_t**

24 **Summary**

25 Lookup published data.

1

Format

PMIx v1.0

C

2

```

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_cbfunc_t cbfunc,
    void *cbdata)

```

C

9

IN proc

`pmix_proc_t` structure of the process seeking the data (handle)

10

11

IN keys

(array of strings)

12

13

IN info

Array of info structures (array of handles)

14

15

IN ninfo

Number of elements in the *info* array (integer)

16

17

IN cbfunc

Callback function `pmix_lookup_cbfunc_t` (function reference)

18

19

IN cbdata

Data to be passed to the callback function (memory reference)

20

21

Returns one of the following:

22

- **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.

23

24

25

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

26

27

- 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

28

Required Attributes

29

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:

30

31

PMIX_USERID "pmix.euid" (uint32_t)

Effective user id.

32

33

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group id.

34

1 Host environments that implement this entry point are required to support the following attributes:

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

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

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

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



Optional Attributes



7 The following attributes are optional for host environments that support this operation:

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

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



12 Description

13 Lookup published data. The host server will be passed a **NULL**-terminated array of string keys
14 identifying the data being requested.

15 The array of *info* structs is used to pass user-requested options to the server. The default data range
16 is left to the host environment, but expected to be **PMIX_SESSION**. This can include a wait flag to
17 indicate that the server should wait for all data to become available before executing the callback
18 function, or should immediately callback with whatever data is available. In addition, a timeout can
19 be specified on the wait to preclude an indefinite wait for data that may never be published.



Advice to PMIx server hosts



20 The **PMIX_USERID** and **PMIX_GRPID** of the requesting process will be provided to support
21 authorization-based access to published information. The host environment is not required to
22 guarantee support for any specific range - i.e., the environment does not need to return an error if
23 the data store doesn't support a specified range so long as it is covered by some internally defined
24 range.

25 10.2.9 pmix_server_unpublish_fn_t

26 Summary

27 Delete data from the data store.

1

Format

PMIx v1.0

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.

1 Host environments that implement this entry point are required to support the following attributes:

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

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

4  **Optional Attributes** 

4 The following attributes are optional for host environments that support this operation:

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

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

9 **Description**

10 Delete data from the data store. The host server will be passed a **NULL**-terminated array of string
11 keys, plus potential directives such as the data range within which the keys should be deleted. The
12 default data range is left to the host environment, but expected to be **PMIX_SESSION**. The
13 callback is to be executed upon completion of the delete procedure.

14  **Advice to PMIx server hosts** 

14 The **PMIX_USERID** and **PMIX_GRPID** of the requesting process will be provided to support
15 authorization-based access to published information. The host environment is not required to
16 guarantee support for any specific range - i.e., the environment does not need to return an error if
17 the data store doesn't support a specified range so long as it is covered by some internally defined
18 range.

19 **10.2.10 pmix_server_spawn_fn_t**

20 **Summary**

21 Spawn a set of applications/processes as per the **PMIx_Spawn** API.

1

Format

PMIx v1.0

C

2

```

typedef pmix_status_t (*pmix_server_spawn_fn_t) (
    const pmix_proc_t *proc,
    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

10

IN `proc`

`pmix_proc_t` structure of the process making the request (handle)

11

12

IN `job_info`

Array of info structures (array of handles)

13

14

IN `ninfo`

Number of elements in the `jobinfo` array (integer)

15

16

IN `apps`

Array of `pmix_app_t` structures (array of handles)

17

18

IN `napps`

Number of elements in the `apps` array (integer)

19

20

IN `cbfunc`

Callback function `pmix_spawn_cbfunc_t` (function reference)

21

22

IN `cbdata`

Data to be passed to the callback function (memory reference)

23

24

Returns one of the following:

25

- **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.

26

27

28

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the `cbfunc` will *not* be called

29

30

- 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

31

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 Host environments that provide this module entry point are required to pass the **PMIX_SPAWNED**
4 and **PMIX_PARENT_ID** attributes to all PMIx servers launching new child processes so those
5 values can be returned to clients upon connection to the PMIx server. In addition, they are required
6 to support the following attributes when present in either the *job_info* or the *info* array of an
7 element of the *apps* array:

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

9 Working directory for spawned processes.

10 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

11 Set the application's current working directory to the session working directory assigned by
12 the RM - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for
13 the rank to discover the session working directory assigned to the provided namespace

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

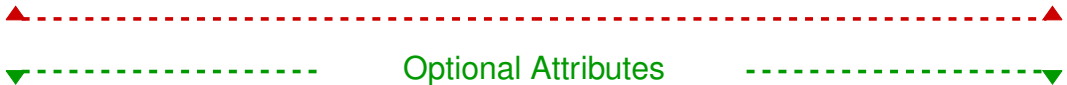
15 Prefix to use for starting spawned processes.

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

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

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

19 Hostfile to use for spawned processes.



20 The following attributes are optional for host environments that support this operation:

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

22 Hostfile listing hosts to add to existing allocation.

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

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

25 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

26 Preload binaries onto nodes.

27 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

28 Comma-delimited list of files to pre-position on nodes.

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

30 Name of personality to use.

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

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

1 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)
 2 Display process mapping upon spawn.

3 **PMIX_PPR** "pmix.ppr" (char*)
 4 Number of processes to spawn on each identified resource.

5 **PMIX_MAPBY** "pmix.mapby" (char*)
 6 Process mapping policy - when accessed using **PMIx_Get** , use the
 7 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
 8 provided namespace

9 **PMIX_RANKBY** "pmix.rankby" (char*)
 10 Process ranking policy - when accessed using **PMIx_Get** , use the
 11 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the
 12 provided namespace

13 **PMIX_BINDTO** "pmix.bindto" (char*)
 14 Process binding policy - when accessed using **PMIx_Get** , use the
 15 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
 16 provided namespace

17 **PMIX_NON_PMI** "pmix.nonpmi" (bool)
 18 Spawned processes will not call **PMIx_Init** .

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

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

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

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

27 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)
 28 Spawned application consists of debugger daemons.

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

31 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
 32 Timestamp output from applications.

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

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

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

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

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

9 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)
10 Do not oversubscribe the cpus.

11 **PMIX_REPORT_BINDINGS** "pmix.repbinding" (bool)
12 Report bindings of the individual processes.

13 **PMIX_CPU_LIST** "pmix.cpulist" (char*)
14 List of cpus to use for this job - when accessed using **PMIx_Get** , use the
15 **PMIX_RANK_WILDCARD** value for the rank to discover the cpu list used for the provided
16 namespace

17 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)
18 Application supports recoverable operations.

19 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)
20 Application is continuous, all failed processes should be immediately restarted.

21 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)
22 Maximum number of times to restart a job - when accessed using **PMIx_Get** , use the
23 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
24 namespace

25 **PMIX_TIMEOUT** "pmix.timeout" (int)
26 Time in seconds before the specified operation should time out (0 indicating infinite) in
27 error. The timeout parameter can help avoid “hangs” due to programming errors that prevent
28 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.

10.2.11 pmix_server_connect_fn_t

Summary

Record the specified processes as *connected*.

Format

PMIx v1.0

C

```
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)
```

C

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:

- 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 host *must not* invoke the callback function
3 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 PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

9 The following attributes are optional for host environments that support this operation:

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

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

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

15 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
16 impose any requirements on a host environment’s collective algorithms. Thus, the
17 acceptable values for this attribute will be environment-dependent - users are encouraged to
18 check their host environment for supported values.

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

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

1 **Description**

2 Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The
3 callback is to be executed once every daemon hosting at least one participant has called the host
4 server's `pmix_server_connect_fn_t` function, *and* the host environment has completed any
5 supporting operations required to meet the terms of the PMIx definition of *connected* processes.

▼————— Advice to PMIx library implementers —————▼

6 The PMIx server library is required to aggregate participation by local clients, passing the request
7 to the host environment once all local participants have executed the API.



▼————— Advice to PMIx server hosts —————▼

8 The host will receive a single call for each collective operation. It is the responsibility of the host to
9 identify the nodes containing participating processes, execute the collective across all participating
10 nodes, and notify the local PMIx server library upon completion of the global collective.



11 **10.2.12 pmix_server_disconnect_fn_t**

12 **Summary**

13 Disconnect a previously connected set of processes.

1

Format

PMIx v1.0

C

2

```

typedef pmix_status_t (*pmix_server_disconnect_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)

```

3

4

5

6

7

8

C

9

IN procs

Array of [pmix_proc_t](#) structures identifying participants (array of handles)

10

11

IN nprocs

Number of elements in the *procs* array (integer)

12

13

IN info

Array of info structures (array of handles)

14

15

IN ninfo

Number of elements in the *info* array (integer)

16

17

IN cbfunc

Callback function [pmix_op_cbfunc_t](#) (function reference)

18

19

IN cbdata

Data to be passed to the callback function (memory reference)

20

21

Returns one of the following:

22

- [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.

23

24

25

- [PMIX_OPERATION_SUCCEEDED](#) , indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

26

27

- 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

28

Required Attributes

29

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.

10.2.13 pmix_server_register_events_fn_t

Summary

Register to receive notifications for the specified events.

1

Format

PMIx v1.0

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)

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.



1

Description

2

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.

3

Advice to PMIx library implementers

4

The PMIx server library must track all client registrations for subsequent notification. This module function shall only be called when:

5

6

- the client has requested notification of an environmental code (i.e., a PMIx code in the range between `PMIX_ERR_SYS_BASE` and `PMIX_ERR_SYS_OTHER`, inclusive) or a code that lies outside the defined PMIx range of constants; and

7

8

9

- 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

10



Advice to PMIx server hosts

11

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).

12

13

14



15 10.2.14 pmix_server_deregister_events_fn_t

16

Summary

17

Deregister to receive notifications for the specified events.

1

Format

PMIx v1.0

C

2

```

typedef pmix_status_t (*pmix_server_deregister_events_fn_t) (
3         pmix_status_t *codes,
4         size_t ncodes,
5         pmix_op_cbfunc_t cbfunc,
6         void *cbdata)

```

C

7

IN codes

Array of `pmix_status_t` values (array of handles)

8

9

IN ncodes

Number of elements in the `codes` array (integer)

10

11

IN cbfunc

Callback function `pmix_op_cbfunc_t` (function reference)

12

13

IN cbdata

Data to be passed to the callback function (memory reference)

14

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

18

19

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the `cbfunc` will *not* be called

20

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

22

23

Description

24

Deregister to receive notifications for the specified events to which the PMIx server has previously registered.

25

Advice to PMIx library implementers

26

The PMIx server library must track all client registrations. This module function shall only be called when:

27

28

- the library is deregistering environmental codes (i.e., a PMIx codes in the range between **PMIX_ERR_SYS_BASE** and **PMIX_ERR_SYS_OTHER**, inclusive) or codes that lies outside the defined PMIx range of constants; and

29

30

31

- no client (including the server library itself) remains registered for notifications on any included code - i.e., a code should be included in this call only when no registered notifications against it remain.

32

33

1 10.2.15 pmix_server_notify_event_fn_t

2 Summary

3 Notify the specified processes of an event.

4 Format

PMIx v2.0

```
typedef pmix_status_t (*pmix_server_notify_event_fn_t)(pmix_status_t code,  
const pmix_proc_t *source,  
pmix_data_range_t range,  
pmix_info_t info[],  
size_t ninfo,  
pmix_op_cbfunc_t cbfunc,  
void *cbdata);
```

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

- 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.

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

10.2.16 pmix_server_listener_fn_t

Summary

Register a socket the host server can monitor for connection requests.

1 **Format**

PMIx v1.0

C

```
2 typedef pmix_status_t (*pmix_server_listener_fn_t) (  
3     int listening_sd,  
4     pmix_connection_cbfunc_t cbfunc,  
5     void *cbdata)
```

C

- 6 **IN** `incoming_sd`
7 (integer)
- 8 **IN** `cbfunc`
9 Callback function `pmix_connection_cbfunc_t` (function reference)
- 10 **IN** `cbdata`
11 (memory reference)

12 Returns **PMIX_SUCCESS** indicating that the request is accepted, or a negative value
13 corresponding to a PMIx error constant indicating that the request has been rejected.

14 **Description**

15 Register a socket the host environment can monitor for connection requests, harvest them, and then
16 call the PMIx server library's internal callback function for further processing. A listener thread is
17 essential to efficiently harvesting connection requests from large numbers of local clients such as
18 occur when running on large SMPs. The host server listener is required to call `accept` on the
19 incoming connection request, and then pass the resulting socket to the provided `cbfunc`. A **NULL**
20 for this function will cause the internal PMIx server to spawn its own listener thread.

21 **10.2.17 pmix_server_query_fn_t**

22 **Summary**

23 Query information from the resource manager.

24 **Format**

PMIx v2.0

C

```
25 typedef pmix_status_t (*pmix_server_query_fn_t) (  
26     pmix_proc_t *proct,  
27     pmix_query_t *queries, size_t nqueries,  
28     pmix_info_cbfunc_t cbfunc,  
29     void *cbdata)
```


1 **IN** `proct`
 2 `pmix_proc_t` structure of the requesting process (handle)
 3 **IN** `queries`
 4 Array of `pmix_query_t` structures (array of handles)
 5 **IN** `nqueries`
 6 Number of elements in the `queries` array (integer)
 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 by the host environment - result
 13 will be returned in the provided `cbfunc`. Note that the host *must not* invoke the callback function
 14 prior to returning 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 • a PMIx error constant indicating either an error in the input or that the request was immediately
 18 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.
 20 In addition, the following attributes are required to be included in the passed `info` array:

21 **PMIX_USERID** "`pmix.euid`" (`uint32_t`)
 22 Effective user id.

23 **PMIX_GRPID** "`pmix.egid`" (`uint32_t`)
 24 Effective group id.

Optional Attributes

25 The following attributes are optional for host environments that support this operation:

26 **PMIX_QUERY_NAMESPACES** "`pmix.qry.ns`" (`char*`)
 27 Request a comma-delimited list of active namespaces.

28 **PMIX_QUERY_JOB_STATUS** "`pmix.qry.jst`" (`pmix_status_t`)
 29 Status of a specified, currently executing job.

30 **PMIX_QUERY_QUEUE_LIST** "`pmix.qry qlst`" (`char*`)
 31 Request a comma-delimited list of scheduler queues.

32 **PMIX_QUERY_QUEUE_STATUS** "`pmix.qry.qst`" (TBD)

1 Status of a specified scheduler queue.

2 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)

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

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

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

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

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

11 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)

12 Return a comma-delimited list of supported debug attributes.

13 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)

14 Return information on memory usage for the processes indicated in the qualifiers.

15 **PMIX_QUERY_LOCAL_ONLY** "pmix.qry.local" (bool)

16 Constrain the query to local information only.

17 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)

18 Report average values.

19 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)

20 Report minimum and maximum values.

21 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)

22 String identifier of the allocation whose status is being requested.

23 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)

24 Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.
25



26 Description

27 Query information from the host environment. The query will include the namespace/rank of the
28 process that is requesting the info, an array of **pmix_query_t** describing the request, and a
29 callback function/data for the return.



Advice to PMIx library implementers

30 The PMIx server library should not block in this function as the host environment may, depending
31 upon the information being requested, require significant time to respond.



1 10.2.18 pmix_server_tool_connection_fn_t

2 Summary

3 Register that a tool has connected to the server.

4 Format

PMIx v2.0

```
▼────────────────────────────────────────── C ───────────────────────────────────────────▼  
5 typedef void (*pmix_server_tool_connection_fn_t) (  
6             pmix_info_t info[], size_t ninfo,  
7             pmix_tool_connection_cbfnc_t cbfunc,  
8             void *cbdata)  
▲────────────────────────────────────────── C ───────────────────────────────────────────▲
```

9 **IN info**
10 Array of [pmix_info_t](#) structures (array of handles)
11 **IN ninfo**
12 Number of elements in the *info* array (integer)
13 **IN cbfunc**
14 Callback function [pmix_tool_connection_cbfnc_t](#) (function reference)
15 **IN cbdata**
16 Data to be passed to the callback function (memory reference)

▼----- Required Attributes -----▼

17 PMIx libraries are required to pass the following attributes in the *info* array:

18 **PMIX_USERID** "pmix.euid" (uint32_t)
19 Effective user id.

20 **PMIX_GRPID** "pmix.egid" (uint32_t)
21 Effective group id.



▼----- Optional Attributes -----▼

22 The following attributes are optional for host environments that support this operation:

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

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

27 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (bool)
28 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.

10.2.19 pmix_server_log_fn_t

Summary

Log data on behalf of a client.

Format

PMIx v2.0

```
typedef void (*pmix_server_log_fn_t) (  
    const pmix_proc_t *client,  
    const pmix_info_t data[], size_t ndata,  
    const pmix_info_t directives[], size_t ndirs,  
    pmix_op_cbfunc_t cbfunc, void *cbdata)
```

IN client
 `pmix_proc_t` structure (handle)

IN data
 Array of info structures (array of handles)

IN ndata
 Number of elements in the *data* array (integer)

IN directives
 Array of info structures (array of handles)

IN ndirs
 Number of elements in the *directives* array (integer)

IN cbfunc
 Callback function `pmix_op_cbfunc_t` (function reference)

IN cbdata
 Data to be passed to the callback function (memory reference)

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.

Host environments that provide this module entry point 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.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_LOG_MSG "pmix.log.msg" (pmix_byte_object_t)

Message blob to be sent somewhere.

PMIX_LOG_EMAIL "pmix.log.email" (pmix_data_array_t)

Log via email based on **pmix_info_t** containing directives.

PMIX_LOG_EMAIL_ADDR "pmix.log.emaddr" (char*)

Comma-delimited list of email addresses that are to receive the message.

PMIX_LOG_EMAIL_SUBJECT "pmix.log.emsub" (char*)

Subject line for email.

PMIX_LOG_EMAIL_MSG "pmix.log.emmsg" (char*)

Message to be included in email.

Description

Log data on behalf of a client. This function is *not* intended for output of computational results, but rather for reporting status and error messages. The host *must not* execute the callback function prior to returning from the API.

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.

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. If not given as part of an `pmix_info_t` struct that identifies the involved nodes, then the description will be applied across all nodes in the requestor's allocation.

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

Name of the network.

1 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)
2 Mbits/sec.
3 **PMIX_ALLOC_NETWORK_QOS** "pmix.alloc.netqos" (char*)
4 Quality of service level.



5 Description

6 Request new allocation or modifications to an existing allocation on behalf of a client. Several
7 broad categories are envisioned, including the ability to:

- 8 • Request allocation of additional resources, including memory, bandwidth, and compute for an
9 existing allocation. Any additional allocated resources will be considered as part of the current
10 allocation, and thus will be released at the same time.
- 11 • Request a new allocation of resources. Note that the new allocation will be disjoint from (i.e., not
12 affiliated with) the allocation of the requestor - thus the termination of one allocation will not
13 impact the other.
- 14 • Extend the reservation on currently allocated resources, subject to scheduling availability and
15 priorities.
- 16 • Return no-longer-required resources to the scheduler. This includes the *loan* of resources back to
17 the scheduler with a promise to return them upon subsequent request.

18 The callback function provides a *status* to indicate whether or not the request was granted, and to
19 provide some information as to the reason for any denial in the `pmix_info_cbfunc_t` array of
20 `pmix_info_t` structures.

21 10.2.21 pmix_server_job_control_fn_t

22 Summary

23 Execute a job control action on behalf of a client.

24 Format

PMIx v2.0

C

```
25 typedef pmix_status_t (*pmix_server_job_control_fn_t) (  
26     const pmix_proc_t *requestor,  
27     const pmix_proc_t targets[], size_t ntargets,  
28     const pmix_info_t directives[], size_t ndirs,  
29     pmix_info_cbfunc_t cbfunc, void *cbdata)
```


1 **IN requestor**
 2 `pmix_proc_t` structure of requesting process (handle)
 3 **IN targets**
 4 Array of proc structures (array of handles)
 5 **IN ntargets**
 6 Number of elements in the *targets* array (integer)
 7 **IN directives**
 8 Array of info structures (array of handles)
 9 **IN ndirs**
 10 Number of elements in the *info* array (integer)
 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

Required Attributes

23 PMIx libraries are required to pass any provided attributes to the host environment for processing.
 24 In addition, the following attributes are required to be included in the passed *info* array:

25 **PMIX_USERID** "pmix.euid" (`uint32_t`)
 26 Effective user id.

27 **PMIX_GRPID** "pmix.egid" (`uint32_t`)
 28 Effective group id.

29 Host environments that provide this module entry point 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.

33 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (`bool`)
 34 Pause the specified processes.

35 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (`bool`)

1 Resume (“un-pause”) the specified processes.
2 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
3 Forcibly terminate the specified processes and cleanup.
4 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
5 Send given signal to specified processes.
6 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
7 Politely terminate the specified processes.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

8 The following attributes are optional for host environments that support this operation:
9 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
10 Cancel the specified request (**NULL** implies cancel all requests from this requestor).
11 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
12 Restart the specified processes using the given checkpoint ID.
13 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
14 Checkpoint the specified processes and assign the given ID to it.
15 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
16 Use event notification to trigger a process checkpoint.
17 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
18 Use the given signal to trigger a process checkpoint.
19 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
20 Time in seconds to wait for a checkpoint to complete.
21 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
22 "pmix.jctrl.ckmethod" (pmix_data_array_t)
23 Array of **pmix_info_t** declaring each method and value supported by this application.
24 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
25 Regular expression identifying nodes that are to be provisioned.
26 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
27 Name of the image that is to be provisioned.
28 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
29 Indicate that the job can be pre-empted.



Description

Execute a job control action on behalf of a client. 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_WILDARD** 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_cbfunc_t** array of **pmix_info_t** structures.

10.2.22 pmix_server_monitor_fn_t

Summary

Request that a client be monitored for activity.

Format

PMIx v2.0

C

```
/* Request that a client be monitored for activity */
typedef pmix_status_t (*pmix_server_monitor_fn_t) (
    const pmix_proc_t *requestor,
    const pmix_info_t *monitor, pmix_status_t error,
    const pmix_info_t directives[], size_t ndirs,
    pmix_info_cbfunc_t cbfunc, void *cbdata);
```

C

- IN requestor**
pmix_proc_t structure of requesting process (handle)
- IN monitor**
pmix_info_t identifying the type of monitor being requested (handle)
- IN error**
Status code to use in generating event if alarm triggers (integer)
- IN directives**
Array of info structures (array of handles)
- IN ndirs**
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)

1 Returns one of the following:

- 2 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
3 will be returned in the provided *cbfunc*. Note that the host *must not* invoke the callback function
4 prior to returning from the API.
- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will *not* be called
- 7 • a PMIx error constant indicating either an error in the input or that the request was immediately
8 processed and failed - the *cbfunc* will *not* be called

9 This entry point is only called for monitoring requests that are not directly supported by the PMIx
10 server library itself.

▼----- Required Attributes -----▼

11 If supported by the PMIx server library, then the library must not pass any supported attributes to
12 the host environment. All attributes not directly supported by the server library must be passed to
13 the host environment if it provides this module entry. In addition, the following attributes are
14 required to be included in the passed *info* array:

15 **PMIX_USERID** "pmix.euid" (uint32_t)
16 Effective user id.

17 **PMIX_GRPID** "pmix.egid" (uint32_t)
18 Effective group id.

19 Host environments are not required to support any specific monitoring attributes.



▼----- Optional Attributes -----▼

20 The following attributes may be implemented by a host environment.

21 **PMIX_MONITOR_ID** "pmix.monitor.id" (char*)
22 Provide a string identifier for this request.

23 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)
24 Identifier to be canceled (**NULL** means cancel all monitoring for this process).

25 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)
26 The application desires to control the response to a monitoring event.

27 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)
28 Register to have the PMIx server monitor the requestor for heartbeats.

29 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)
30 Time in seconds before declaring heartbeat missed.

31 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrop" (uint32_t)

1 Number of heartbeats that can be missed before generating the event.
2 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)
3 Register to monitor file for signs of life.
4 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)
5 Monitor size of given file is growing to determine if the application is running.
6 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)
7 Monitor time since last access of given file to determine if the application is running.
8 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
9 Monitor time since last modified of given file to determine if the application is running.
10 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
11 Time in seconds between checking the file.
12 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)
13 Number of file checks that can be missed before generating the event.



14 **Description**

15 Request that a client be monitored for activity.



Advice to PMIx server hosts

16 If this module entry is provided and called by the PMIx server library, then the host environment
17 must either provide the requested services or return **PMIX_ERR_NOT_SUPPORTED** to the
18 provided *cbfunc*.



APPENDIX A

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 **A.1 Version 2.0**

6 The following list includes some of the active participants in the PMIx v2 standardization process.

- 7 • Ralph H. Castain, Annapurna Dasari, Christopher A. Holguin, Andrew Friedley, Michael Klemm
8 and Terry Wilmarth
- 9 • Joshua Hursey, David Solt, Alexander Eichenberger, Geoff Paulsen, and Sameh Sharkawi
- 10 • Aurelien Bouteiller and George Bosilca
- 11 • Artem Polyakov, Igor Ivanov and Boris Karasev
- 12 • Gilles Gouaillardet
- 13 • Michael A Raymond and Jim Stoffel
- 14 • Dirk Schubert
- 15 • Moe Jette
- 16 • Takahiro Kawashima and Shinji Sumimoto
- 17 • Howard Pritchard
- 18 • David Beer
- 19 • Brice Goglin
- 20 • Geoffroy Vallee, Swen Boehm, Thomas Naughton and David Bernholdt
- 21 • Adam Moody and Martin Schulz
- 22 • Ryan Grant and Stephen Olivier
- 23 • Michael Karo

1 The following institutions supported this effort through time and travel support for the people listed
2 above.

- 3 • Intel Corporation
- 4 • IBM, Inc.
- 5 • University of Tennessee, Knoxville
- 6 • The Exascale Computing Project, an initiative of the US Department of Energy
- 7 • National Science Foundation
- 8 • Mellanox, Inc.
- 9 • Research Organization for Information Science and Technology
- 10 • HPE Co.
- 11 • Allinea (ARM)
- 12 • SchedMD, Inc.
- 13 • Fujitsu Limited
- 14 • Los Alamos National Laboratory
- 15 • Adaptive Solutions, Inc.
- 16 • INRIA
- 17 • Oak Ridge National Laboratory
- 18 • Lawrence Livermore National Laboratory
- 19 • Sandia National Laboratory
- 20 • Altair

21 **A.2 Version 1.0**

22 The following list includes some of the active participants in the PMIx v1 standardization process.

- 23 • Ralph H. Castain, Annapurna Dasari and Christopher A. Holguin
- 24 • Joshua Hursey and David Solt
- 25 • Aurelien Bouteiller and George Bosilca
- 26 • Artem Polyakov, Elena Shipunova, Igor Ivanov, and Joshua Ladd
- 27 • Gilles Gouaillardet
- 28 • Gary Brown

1
2
3
4
5
6
7
8
9
10

- Moe Jette

The following institutions supported this effort through time and travel support for the people listed above.

- Intel Corporation
- IBM, Inc.
- University of Tennessee, Knoxville
- Mellanox, Inc.
- Research Organization for Information Science and Technology
- Adaptive Solutions, Inc.
- 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, [10](#), [57](#), [58](#), [99](#), [107](#), [155](#), [194](#),
[197](#)
 - Defintion, [12](#)
- host environment
 - Defintion, [13](#)
- job, [10](#), [57](#), [58](#), [99](#), [107](#), [109](#), [155](#), [190](#), [193](#),
[194](#), [196](#), [197](#)
 - Defintion, [12](#)
- namespace
 - Defintion, [12](#)
- PMIx_Abort, [8](#), [25](#), [130](#), [213](#), [214](#)
 - Defintion, [129](#)
- PMIX_ADD_HOST, [132](#), [136](#), [226](#)
 - Defintion, [64](#)
- PMIX_ADD_HOSTFILE, [132](#), [136](#), [226](#)
 - Defintion, [64](#)
- PMIX_ALLOC_BANDWIDTH, [157](#), [247](#)
 - Defintion, [69](#)
- PMIX_ALLOC_CPU_LIST, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_DIRECTIVE, [51](#)
- PMIx_Alloc_directive_string, [9](#)
 - Defintion, [85](#)
- pmix_alloc_directive_t, [37](#), [51](#), [85](#), [245](#)
 - Defintion, [37](#)
- PMIX_ALLOC_EXTEND, [37](#)
- PMIX_ALLOC_EXTERNAL, [37](#)
- PMIX_ALLOC_ID, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_MEM_SIZE, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NETWORK, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NETWORK_ID, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NETWORK_QOS, [158](#),
[247](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NEW, [37](#)
- PMIX_ALLOC_NODE_LIST, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NUM_CPU_LIST, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NUM_CPUS, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_NUM_NODES, [157](#), [246](#)
 - Defintion, [69](#)
- PMIX_ALLOC_REAQUIRE, [37](#)
- PMIX_ALLOC_RELEASE, [37](#)
- PMIX_ALLOC_TIME, [157](#), [246](#)
 - Defintion, [70](#)
- PMIX_ALLOCATED_NODELIST, [191](#)
 - Defintion, [57](#)
- PMIx_Allocation_request_nb, [9](#), [69](#), [148](#),
[158](#)
 - Defintion, [156](#)
- PMIX_ANL_MAP
 - Defintion, [62](#)
- PMIX_APP, [51](#)
- PMIX_APP_CONSTRUCT
 - Defintion, [41](#)
- PMIX_APP_CREATE
 - Defintion, [42](#)
- PMIX_APP_DESTRUCT
 - Defintion, [42](#)
- PMIX_APP_FREE
 - Defintion, [42](#)
- PMIX_APP_INFO, [102](#), [105](#), [110](#), [151](#)
 - Defintion, [58](#)

PMIX_APP_INFO_ARRAY, 197
 Defintion, 58
 PMIX_APP_MAP_REGEX
 Defintion, 62
 PMIX_APP_MAP_TYPE
 Defintion, 62
 PMIX_APP_RANK, 190
 Defintion, 56
 PMIX_APP_SIZE, 109, 190, 197
 Defintion, 59
 pmix_app_t, 41, 42, 131, 135, 225
 Defintion, 41
 PMIX_APPLDR, 190, 197
 Defintion, 56
 PMIX_APPNUM, 58, 102, 105, 109, 151,
 190, 197
 Defintion, 56
 PMIX_ARCH
 Defintion, 55
 PMIX_ATTR_UNDEF
 Defintion, 52
 PMIX_AVAIL_PHYS_MEMORY, 191
 Defintion, 59
 PMIX_BINDTO, 133, 137, 191, 227
 Defintion, 65
 PMIX_BOOL, 50
 PMIX_BUFFER, 51
 PMIX_BYTE, 50
 PMIX_BYTE_OBJECT, 51
 PMIX_BYTE_OBJECT_CREATE
 Defintion, 47
 PMIX_BYTE_OBJECT_DESTRUCT
 Defintion, 47
 PMIX_BYTE_OBJECT_FREE
 Defintion, 47
 PMIX_BYTE_OBJECT_LOAD
 Defintion, 48
 pmix_byte_object_t, 46–48, 51
 Defintion, 46
 PMIX_CLIENT_AVG_MEMORY
 Defintion, 59
 PMIX_COLLECT_DATA, 113, 115, 215
 Defintion, 60
 PMIX_COLLECTIVE_ALGO, 10, 113,
 116, 140, 143, 216, 230
 Defintion, 61
 PMIX_COLLECTIVE_ALGO_REQD, 113,
 116, 140, 143, 216, 230
 Defintion, 61
 PMIX_COMMAND, 51
 PMIx_Commit, 8, 82, 100, 112, 139, 205,
 218
 Defintion, 112
 PMIX_COMPRESSED_STRING, 51
 PMIx_Connect, 8, 10, 19, 134, 139, 141,
 142, 144, 145
 Defintion, 139
 PMIX_CONNECT_MAX_RETRIES, 92
 Defintion, 53
 PMIx_Connect_nb, 8, 142
 Defintion, 142
 PMIX_CONNECT_RETRY_DELAY, 92
 Defintion, 53
 PMIX_CONNECT_SYSTEM_FIRST, 92,
 94
 Defintion, 53
 PMIX_CONNECT_TO_SYSTEM, 92, 94
 Defintion, 53
 pmix_connection_cbfunc_t, 239
 Defintion, 82
 PMIX_COSPAWN_APP
 Defintion, 65
 PMIX_CPU_LIST, 134, 138, 228
 Defintion, 66
 PMIX_CPUS_PER_PROC, 133, 138, 228
 Defintion, 66
 PMIX_CPuset
 Defintion, 55
 PMIX_CREDENTIAL
 Defintion, 55
 PMIX_DAEMON_MEMORY
 Defintion, 59
 PMIX_DATA_ARRAY, 51
 pmix_data_array_t, 29, 50, 51, 67, 153,
 196–198, 241
 Defintion, 50

PMIX_DATA_BUFFER_CONSTRUCT,
 180, 182
 Defintion, [177](#)
 PMIX_DATA_BUFFER_CREATE, [180](#),
 182
 Defintion, [49](#), [176](#)
 PMIX_DATA_BUFFER_DESTRUCT
 Defintion, [49](#), [177](#)
 PMIX_DATA_BUFFER_LOAD
 Defintion, [178](#)
 PMIX_DATA_BUFFER_RELEASE
 Defintion, [49](#), [177](#)
 pmix_data_buffer_t, [48](#), [49](#), [176–181](#), [185](#)
 Defintion, [48](#)
 PMIX_DATA_BUFFER_UNLOAD
 Defintion, [178](#)
 PMIx_Data_copy, [9](#)
 Defintion, [183](#)
 PMIx_Data_copy_payload, [9](#)
 Defintion, [184](#)
 PMIx_Data_pack, [9](#), [180](#)
 Defintion, [179](#)
 PMIx_Data_print, [9](#)
 Defintion, [183](#)
 PMIX_DATA_RANGE, [51](#)
 PMIx_Data_range_string, [9](#)
 Defintion, [84](#)
 pmix_data_range_t, [28](#), [51](#), [84](#), [174](#), [237](#)
 Defintion, [28](#)
 PMIX_DATA_SCOPE, [101](#), [105](#)
 Defintion, [61](#)
 PMIX_DATA_TYPE, [51](#)
 PMIX_DATA_TYPE_MAX, [51](#)
 PMIx_Data_type_string, [9](#)
 Defintion, [85](#)
 pmix_data_type_t, [31](#), [34](#), [40](#), [50](#), [85](#), [180](#),
 181, [183](#), [184](#)
 Defintion, [50](#)
 PMIx_Data_unpack, [9](#)
 Defintion, [181](#)
 PMIX_DEBUG_JOB
 Defintion, [68](#)
 PMIX_DEBUG_STOP_IN_INIT
 Defintion, [68](#)
 PMIX_DEBUG_STOP_ON_EXEC
 Defintion, [68](#)
 PMIX_DEBUG_WAIT_FOR_NOTIFY
 Defintion, [68](#)
 PMIX_DEBUG_WAITING_FOR_NOTIFY
 Defintion, [68](#)
 PMIX_DEBUGGER_DAEMONS, [133](#), [137](#),
 227
 Defintion, [65](#)
 PMIx_Deregister_event_handler, [9](#)
 Defintion, [172](#)
 PMIx_Disconnect, [8](#), [10](#), [19](#), [142](#), [145–147](#)
 Defintion, [144](#)
 PMIx_Disconnect_nb, [8](#), [147](#)
 Defintion, [146](#)
 PMIX_DISPLAY_MAP, [132](#), [137](#), [227](#)
 Defintion, [65](#)
 pmix_dmodex_response_fn_t, [205](#)
 Defintion, [81](#)
 PMIX_DOUBLE, [51](#)
 PMIX_DSTPATH
 Defintion, [54](#)
 PMIX_EMBED_BARRIER, [90](#)
 Defintion, [61](#)
 PMIX_ERR_BAD_PARAM, [18](#)
 PMIX_ERR_COMM_FAILURE, [19](#)
 PMIX_ERR_DATA_VALUE_NOT_FOUND,
 18
 PMIX_ERR_DEBUGGER_RELEASE, [18](#)
 PMIX_ERR_EVENT_REGISTRATION, [19](#)
 PMIX_ERR_HANDSHAKE_FAILED, [18](#)
 PMIX_ERR_IN_ERRNO, [18](#)
 PMIX_ERR_INIT, [18](#)
 PMIX_ERR_INVALID_ARG, [18](#)
 PMIX_ERR_INVALID_ARGS, [19](#)
 PMIX_ERR_INVALID_CRED, [18](#)
 PMIX_ERR_INVALID_KEY, [18](#)
 PMIX_ERR_INVALID_KEY_LENGTH, [18](#)
 PMIX_ERR_INVALID_KEYVALP, [19](#)
 PMIX_ERR_INVALID_LENGTH, [19](#)
 PMIX_ERR_INVALID_NAMESPACE, [19](#)
 PMIX_ERR_INVALID_NUM_ARGS, [19](#)

PMIX_ERR_INVALID_NUM_PARSED, 19
 PMIX_ERR_INVALID_OPERATION, 20
 PMIX_ERR_INVALID_SIZE, 19
 PMIX_ERR_INVALID_TERMINATION, 19
 PMIX_ERR_INVALID_VAL, 18
 PMIX_ERR_INVALID_VAL_LENGTH, 19
 PMIX_ERR_JOB_TERMINATED, 19
 PMIX_ERR_LOST_CONNECTION_TO_CLIENT, 19
 PMIX_ERR_LOST_CONNECTION_TO_SERVER, 19
 PMIX_ERR_LOST_PEER_CONNECTION, 19
 PMIX_ERR_NO_PERMISSIONS, 18
 PMIX_ERR_NODE_DOWN, 20
 PMIX_ERR_NODE_OFFLINE, 20
 PMIX_ERR_NOMEM, 18
 PMIX_ERR_NOT_FOUND, 19
 PMIX_ERR_NOT_IMPLEMENTED, 19
 PMIX_ERR_NOT_SUPPORTED, 19
 PMIX_ERR_OUT_OF_RESOURCE, 18
 PMIX_ERR_PACK_FAILURE, 18
 PMIX_ERR_PACK_MISMATCH, 18
 PMIX_ERR_PROC_ABORTED, 18
 PMIX_ERR_PROC_ABORTING, 18
 PMIX_ERR_PROC_CHECKPOINT, 18
 PMIX_ERR_PROC_ENTRY_NOT_FOUND, 18
 PMIX_ERR_PROC_MIGRATE, 18
 PMIX_ERR_PROC_REQUESTED_ABORT, 18
 PMIX_ERR_PROC_RESTART, 18
 PMIX_ERR_READY_FOR_HANDSHAKE, 18
 PMIX_ERR_RESOURCE_BUSY, 18
 PMIX_ERR_SERVER_FAILED_REQUEST, 18
 PMIX_ERR_SERVER_NOT_AVAIL, 19
 PMIX_ERR_SILENT, 18
 PMIX_ERR_TIMEOUT, 18
 PMIX_ERR_TYPE_MISMATCH, 18
 PMIX_ERR_UNKNOWN_DATA_TYPE, 18
 PMIX_ERR_UNPACK_FAILURE, 18
 PMIX_ERR_UNPACK_INADEQUATE_SPACE, 18
 PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER, 19
 PMIX_ERR_UNREACH, 18
 PMIX_ERR_UPDATE_ENDPOINTS, 19
 PMIX_ERR_WOULD_BLOCK, 18
 PMIX_ERROR, 18
 PMIX_ERROR_GROUP_ABORT
 Defintion, 62
 PMIX_ERROR_GROUP_COMM
 Defintion, 62
 PMIX_ERROR_GROUP_GENERAL
 Defintion, 63
 PMIX_ERROR_GROUP_LOCAL
 Defintion, 63
 PMIX_ERROR_GROUP_MIGRATE
 Defintion, 62
 PMIX_ERROR_GROUP_NODE
 Defintion, 62
 PMIX_ERROR_GROUP_RESOURCE
 Defintion, 62
 PMIX_ERROR_GROUP_SPAWN
 Defintion, 62
 PMIX_ERROR_HANDLER_ID
 Defintion, 63
 PMIX_ERROR_NAME
 Defintion, 62
 PMIx_Error_string, 9
 Defintion, 84
 PMIX_EVENT_ACTION_COMPLETE, 20
 PMIX_EVENT_ACTION_DEFERRED, 20
 PMIX_EVENT_ACTION_TIMEOUT, 171
 Defintion, 64
 PMIX_EVENT_AFFECTED_PROC, 171, 174
 Defintion, 63
 PMIX_EVENT_AFFECTED_PROCS, 171, 175
 Defintion, 63
 PMIX_EVENT_BASE, 89, 93, 97

Definition, [52](#)
 PMIX_EVENT_CUSTOM_RANGE, [170](#),
 [174](#)
 Definition, [63](#)
 PMIX_EVENT_DO_NOT_CACHE
 Definition, [63](#)
 PMIX_EVENT_HDLR_AFTER, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_APPEND, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_BEFORE, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_FIRST, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_FIRST_IN_CATEGORY,
 [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_LAST, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_LAST_IN_CATEGORY,
 [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_NAME, [170](#)
 Definition, [63](#)
 PMIX_EVENT_HDLR_PREPEND, [170](#)
 Definition, [63](#)
 PMIX_EVENT_NO_ACTION_TAKEN, [20](#)
 PMIX_EVENT_NO_TERMINATION
 Definition, [64](#)
 PMIX_EVENT_NON_DEFAULT, [174](#)
 Definition, [63](#)
 pmix_event_notification_cbfunc_fn_t, [77](#), [79](#)
 Definition, [77](#)
 PMIX_EVENT_PARTIAL_ACTION_TAKEN,
 [20](#)
 PMIX_EVENT_RETURN_OBJECT, [171](#)
 Definition, [63](#)
 PMIX_EVENT_SILENT_TERMINATION,
 [171](#)
 Definition, [63](#)
 PMIX_EVENT_TERMINATE_JOB, [171](#)
 Definition, [64](#)
 PMIX_EVENT_TERMINATE_NODE, [171](#)
 Definition, [64](#)
 PMIX_EVENT_TERMINATE_PROC, [171](#)
 Definition, [64](#)
 PMIX_EVENT_TERMINATE_SESSION,
 [171](#)
 Definition, [64](#)
 PMIX_EVENT_WANT_TERMINATION
 Definition, [64](#)
 pmix_evhdr_reg_cbfunc_t, [76](#), [170](#)
 Definition, [76](#)
 PMIX_EXISTS, [18](#)
 PMIX_EXTERNAL_ERR_BASE, [20](#)
 PMIx_Fence, [3](#), [7](#), [8](#), [11](#), [97](#), [114](#), [116](#), [139](#),
 [142](#), [146](#), [205](#), [214](#), [216](#)
 Definition, [112](#)
 PMIx_Fence_nb, [8](#), [74](#), [116](#), [214](#), [216](#)
 Definition, [114](#)
 PMIx_Finalize, [8](#), [19](#), [25](#), [61](#), [89](#), [90](#), [139](#),
 [212](#), [213](#)
 Definition, [90](#)
 PMIX_FLOAT, [51](#)
 PMIX_FWD_STDERR, [133](#), [137](#), [227](#), [242](#)
 Definition, [65](#)
 PMIX_FWD_STDIN, [133](#), [137](#), [227](#), [242](#)
 Definition, [65](#)
 PMIX_FWD_STDOUT, [133](#), [137](#), [227](#), [242](#)
 Definition, [65](#)
 PMIX_GDS_ACTION_COMPLETE, [20](#)
 PMIX_GDS_MODULE, [89](#), [93](#), [97](#)
 Definition, [55](#)
 PMIx_generate_ppn, [8](#)
 Definition, [187](#)
 PMIx_generate_regex, [8](#), [193](#)
 Definition, [186](#)
 PMIx_Get, [3](#), [8](#), [10](#), [29](#), [52–55](#), [60–62](#),
 [64–71](#), [89](#), [101](#), [103](#), [105–109](#),
 [111](#), [132–134](#), [136–138](#), [154](#), [155](#),
 [189](#), [191](#), [226–228](#)
 Definition, [100](#)
 PMIx_Get_nb, [8](#), [75](#)
 Definition, [103](#)
 PMIx_Get_version, [9](#), [14](#)
 Definition, [87](#)

PMIX_GLOBAL, 28
 PMIX_GLOBAL_RANK, 192
 Defintion, 56
 PMIX_GRPID, 118, 120, 122, 124, 126,
 128, 152, 157, 159, 162, 165,
 219–224, 226, 234, 240, 242, 244,
 246, 248, 251
 Defintion, 54
 PMIx_Heartbeat, 9
 Defintion, 164
 PMIX_HOST, 132, 136, 226
 Defintion, 64
 PMIX_HOSTFILE, 132, 136, 226
 Defintion, 64
 PMIX_HOSTNAME, 58, 102, 105, 111,
 152, 192
 Defintion, 57
 PMIX_HWLOC_SHMEM_ADDR
 Defintion, 60
 PMIX_HWLOC_SHMEM_FILE
 Defintion, 60
 PMIX_HWLOC_SHMEM_SIZE
 Defintion, 60
 PMIX_HWLOC_XML_V1, 191
 Defintion, 60
 PMIX_HWLOC_XML_V2, 191
 Defintion, 60
 PMIX_IMMEDIATE, 101, 105
 Defintion, 60
 PMIX_INDEX_ARGV, 133, 138, 228
 Defintion, 66
 PMIX_INFO, 51
 PMIX_INFO_ARRAY, 51
 pmix_info_array, 32, 33
 Defintion, 32
 pmix_info_cbfunc_t, 72, 75, 150, 156, 159,
 161, 162, 164, 240, 245, 247, 250
 Defintion, 75
 PMIX_INFO_CONSTRUCT
 Defintion, 33
 PMIX_INFO_CREATE
 Defintion, 34
 PMIX_INFO_DESTRUCT
 Defintion, 33
 PMIX_INFO_DIRECTIVES, 51
 PMIx_Info_directives_string, 9
 Defintion, 85
 pmix_info_directives_t, 36, 85
 Defintion, 36
 PMIX_INFO_FREE
 Defintion, 34
 PMIX_INFO_IS_REQUIRED, 36
 Defintion, 37
 PMIX_INFO_LOAD
 Defintion, 34
 PMIX_INFO_REQD, 36
 PMIX_INFO_REQUIRED, 36
 Defintion, 37
 pmix_info_t, 3, 9, 11, 28, 32–37, 58, 68–70,
 76–79, 88, 90, 91, 95, 97, 119,
 123, 154, 156, 157, 160, 161, 164,
 166, 174, 193, 196–198, 237,
 242–244, 246, 247, 249, 250
 Defintion, 32
 PMIX_INFO_TRUE
 Defintion, 35
 PMIX_INFO_XFER, 193
 Defintion, 35
 PMIx_Init, 9, 65, 68, 87, 89, 90, 133, 137,
 211, 227
 Defintion, 87
 PMIx_Initialized, 8
 Defintion, 86
 PMIX_INT, 50
 PMIX_INT16, 50
 PMIX_INT32, 50
 PMIX_INT64, 50
 PMIX_INT8, 50
 PMIX_INTERNAL, 28
 PMIX_JCTRL_CHECKPOINT, 19
 PMIX_JCTRL_CHECKPOINT_COMPLETE,
 19
 PMIX_JCTRL_PREEMPT_ALERT, 19
 PMIX_JOB_CONTINUOUS, 134, 138, 228
 Defintion, 66
 PMIx_Job_control_nb, 9, 70, 148, 158, 193

Definition, [158](#)
 PMIX_JOB_CTRL_CANCEL, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_CHECKPOINT, [160](#),
 [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_CHECKPOINT_EVENT,
 [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_CHECKPOINT_METHOD,
 [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_CHECKPOINT_SIGNAL,
 [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT,
 [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_ID, [159](#), [248](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_KILL, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_PAUSE, [160](#), [248](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_PREEMPTIBLE, [160](#),
 [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_PROVISION, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_PROVISION_IMAGE,
 [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_RESTART, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_RESUME, [160](#), [248](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_SIGNAL, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_CTRL_TERMINATE, [160](#), [249](#)
 Definition, [70](#)
 PMIX_JOB_INFO, [101](#), [105](#), [109](#), [151](#)
 Definition, [57](#)
 PMIX_JOB_INFO_ARRAY, [196](#)
 Definition, [58](#)
 PMIX_JOB_NUM_APPS, [109](#), [191](#), [196](#)
 Definition, [59](#)
 PMIX_JOB_RECOVERABLE, [134](#), [138](#),
 [228](#)
 Definition, [66](#)
 PMIX_JOB_SIZE, [10](#), [103](#), [106](#), [109](#), [189](#),
 [196](#), [197](#)
 Definition, [59](#)
 PMIX_JOB_TERM_STATUS
 Definition, [61](#)
 PMIX_JOBID, [58](#), [102](#), [105](#), [109](#), [151](#), [189](#),
 [196](#)
 Definition, [56](#)
 pmix_key_t, [20](#), [21](#), [99](#), [101](#)
 Definition, [20](#)
 PMIX_KVAL, [51](#)
 PMIX_LOCAL, [27](#)
 PMIX_LOCAL_CPUSSETS, [190](#), [199](#)
 Definition, [57](#)
 PMIX_LOCAL_PEERS, [189](#), [199](#)
 Definition, [57](#)
 PMIX_LOCAL_PROCS, [192](#)
 Definition, [57](#)
 PMIX_LOCAL_RANK, [152](#), [153](#), [190](#)
 Definition, [56](#)
 PMIX_LOCAL_SIZE, [189](#)
 Definition, [59](#)
 PMIX_LOCAL_TOPO
 Definition, [59](#)
 PMIX_LOCALITY
 Definition, [57](#)
 PMIX_LOCALITY_STRING
 Definition, [60](#)
 PMIX_LOCALLDR, [192](#)
 Definition, [56](#)
 PMIX_LOG_EMAIL, [166](#), [244](#)
 Definition, [68](#)
 PMIX_LOG_EMAIL_ADDR, [166](#), [244](#)
 Definition, [68](#)
 PMIX_LOG_EMAIL_MSG, [166](#), [244](#)
 Definition, [68](#)
 PMIX_LOG_EMAIL_SUBJECT, [166](#), [244](#)

Defintion, [68](#)
 PMIX_LOG_MSG, [166](#), [244](#)
 Defintion, [68](#)
 PMIx_Log_nb, [9](#), [67](#), [166](#)
 Defintion, [164](#)
 PMIX_LOG_STDERR, [166](#), [244](#)
 Defintion, [67](#)
 PMIX_LOG_STDOUT, [166](#), [244](#)
 Defintion, [67](#)
 PMIX_LOG_SYSLOG, [166](#), [244](#)
 Defintion, [68](#)
 PMIx_Lookup, [8](#), [38](#), [117](#), [123](#), [125](#)
 Defintion, [121](#)
 pmix_lookup_cbfunc_t, [74](#), [221](#)
 Defintion, [74](#)
 PMIx_Lookup_nb, [74](#), [75](#)
 Defintion, [123](#)
 PMIX_MAP_BLOB
 Defintion, [62](#)
 PMIX_MAPBY, [132](#), [137](#), [191](#), [227](#)
 Defintion, [65](#)
 PMIX_MAPPER, [64](#), [132](#), [136](#), [226](#)
 Defintion, [65](#)
 PMIX_MAX_KEYLEN, [16](#)
 PMIX_MAX_NSLEN, [16](#)
 PMIX_MAX_PROCS, [189](#)
 Defintion, [59](#)
 PMIX_MAX_RESTARTS, [134](#), [138](#), [228](#)
 Defintion, [66](#)
 PMIX_MERGE_STDERR_STDOUT, [133](#),
 [137](#), [227](#)
 Defintion, [66](#)
 PMIX_MODEL_DECLARED, [20](#)
 PMIX_MODEL_LIBRARY_NAME
 Defintion, [54](#)
 PMIX_MODEL_LIBRARY_VERSION
 Defintion, [54](#)
 PMIX_MODEX, [51](#)
 pmix_modex_cbfunc_t, [72](#), [215](#), [217](#)
 Defintion, [72](#)
 PMIX_MODEX_CONSTRUCT
 Defintion, [45](#)
 PMIX_MODEX_CREATE
 Defintion, [45](#)
 pmix_modex_data_t, [44–46](#)
 Defintion, [44](#)
 PMIX_MODEX_DESTRUCT
 Defintion, [45](#)
 PMIX_MODEX_FREE
 Defintion, [46](#)
 PMIX_MONITOR_APP_CONTROL, [163](#),
 [251](#)
 Defintion, [71](#)
 PMIX_MONITOR_CANCEL, [162](#), [251](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE, [163](#), [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE_ACCESS, [163](#),
 [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE_ALERT, [19](#)
 PMIX_MONITOR_FILE_CHECK_TIME,
 [163](#), [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE_DROPS, [163](#), [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE_MODIFY, [163](#),
 [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_FILE_SIZE, [163](#), [252](#)
 Defintion, [71](#)
 PMIX_MONITOR_HEARTBEAT, [163](#), [251](#)
 Defintion, [71](#)
 PMIX_MONITOR_HEARTBEAT_ALERT,
 [19](#)
 PMIX_MONITOR_HEARTBEAT_DROPS,
 [163](#), [251](#)
 Defintion, [71](#)
 PMIX_MONITOR_HEARTBEAT_TIME,
 [163](#), [251](#)
 Defintion, [71](#)
 PMIX_MONITOR_ID, [162](#), [251](#)
 Defintion, [71](#)
 PMIX_NET_TOPO
 Defintion, [59](#)
 PMIX_NO_OVERSUBSCRIBE, [133](#), [138](#),

228
 Definition, **66**
 PMIX_NO_PROCS_ON_HEAD, 133, 138,
 228
 Definition, **66**
 PMIX_NODE_INFO, 102, 105, 111, 151
 Definition, **58**
 PMIX_NODE_INFO_ARRAY, 198
 Definition, **58**
 PMIX_NODE_LIST
 Definition, **57, 59**
 PMIX_NODE_MAP, 57, 189, 190, 196, 197
 Definition, **61**
 PMIX_NODE_RANK, 190
 Definition, **56**
 PMIX_NODE_SIZE, 111, 192
 Definition, **59**
 PMIX_NODEID, 58, 102, 105, 111, 152,
 190
 Definition, **57**
 PMIX_NON_PMI, 133, 137, 227
 Definition, **65**
 pmix_notification_fn_t, 78, 170
 Definition, **78**
 PMIX_NOTIFY_ALLOC_COMPLETE, 19
 PMIX_NOTIFY_COMPLETION
 Definition, **61**
 PMIx_Notify_event, 9
 Definition, **173**
 PMIX_NPROC_OFFSET, 191
 Definition, **56**
 PMIX_NSDIR, 56
 Definition, **56**
 PMIX_NSPACE, 58, 102, 105, 109,
 151–153, 197
 Definition, **56**
 pmix_nspace_t, 21, 24, 73
 Definition, **21**
 PMIX_NUM_NODES, 99, 103, 106, 107,
 109, 110, 196, 197
 Definition, **59**
 PMIX_NUM_SLOTS, 59, 111, 189
 Definition, **59**
 pmix_op_cbfunc_t, 74, 77, 81, 120, 127,
 142, 146, 165, 173, 174, 188,
 201–203, 208, 211, 212, 214, 219,
 223, 229, 232, 234, 236, 237, 243,
 248, 250
 Definition, **74**
 PMIX_OPTIONAL, 101, 104
 Definition, **61**
 PMIX_OUTPUT_TO_FILE, 133, 137, 227
 Definition, **66**
 PMIX_PARENT_ID, 131, 135, 136, 226
 Definition, **57**
 PMIX_PDATA, 51
 PMIX_PDATA_CONSTRUCT
 Definition, **38**
 PMIX_PDATA_CREATE
 Definition, **39**
 PMIX_PDATA_DESTRUCT
 Definition, **38**
 PMIX_PDATA_FREE
 Definition, **39**
 PMIX_PDATA_LOAD
 Definition, **39**
 pmix_pdata_t, 38–40, 74, 75, 123
 Definition, **38**
 PMIX_PDATA_XFER
 Definition, **40**
 PMIX_PERSIST, 51
 PMIX_PERSIST_APP, 28
 PMIX_PERSIST_FIRST_READ, 28
 PMIX_PERSIST_INDEF, 28
 PMIX_PERSIST_PROC, 28
 PMIX_PERSIST_SESSION, 28
 PMIX_PERSISTENCE, 118, 121, 220
 Definition, **61**
 PMIx_Persistence_string, 9
 Definition, **84**
 pmix_persistence_t, 28, 51, 84
 Definition, **28**
 PMIX_PERSONALITY, 132, 136, 226
 Definition, **64**
 PMIX_PID, 50
 PMIX_POINTER, 51

PMIX_PPR, 132, 137, 227
 Defintion, 65
 PMIX_PREFIX, 132, 136, 226
 Defintion, 64
 PMIX_PRELOAD_BIN, 132, 136, 226
 Defintion, 65
 PMIX_PRELOAD_FILES, 132, 136, 226
 Defintion, 65
 PMIX_PROC, 51
 PMIX_PROC_BLOB
 Defintion, 62
 PMIX_PROC_CONSTRUCT, 23
 Defintion, 23, 46, 49
 PMIX_PROC_CREATE
 Defintion, 23
 PMIX_PROC_DATA, 197
 Defintion, 61
 PMIX_PROC_DESTRUCT
 Defintion, 23
 PMIX_PROC_FREE, 149
 Defintion, 24
 PMIX_PROC_INFO, 51
 PMIX_PROC_INFO_CONSTRUCT
 Defintion, 26
 PMIX_PROC_INFO_CREATE
 Defintion, 27
 PMIX_PROC_INFO_DESTRUCT
 Defintion, 26
 PMIX_PROC_INFO_FREE
 Defintion, 27
 pmix_proc_info_t, 25–27, 51, 67, 153, 241
 Defintion, 25
 PMIX_PROC_LOAD
 Defintion, 24
 PMIX_PROC_MAP, 189, 196, 197
 Defintion, 61
 PMIX_PROC_PID
 Defintion, 56
 PMIX_PROC_RANK, 51
 PMIX_PROC_STATE, 51
 PMIX_PROC_STATE_ABORTED, 25
 PMIX_PROC_STATE_ABORTED_BY_SIG,
 25
 PMIX_PROC_STATE_CALLED_ABORT,
 25
 PMIX_PROC_STATE_CANNOT_RESTART,
 25
 PMIX_PROC_STATE_COMM_FAILED,
 25
 PMIX_PROC_STATE_CONNECTED, 25
 PMIX_PROC_STATE_ERROR, 25
 PMIX_PROC_STATE_FAILED_TO_LAUNCH,
 25
 PMIX_PROC_STATE_FAILED_TO_START,
 25
 PMIX_PROC_STATE_KILLED_BY_CMD,
 25
 PMIX_PROC_STATE_LAUNCH_UNDERWAY,
 25
 PMIX_PROC_STATE_MIGRATING, 25
 PMIX_PROC_STATE_PREPPED, 25
 PMIX_PROC_STATE_RESTART, 25
 PMIX_PROC_STATE_RUNNING, 25
 PMIX_PROC_STATE_STATUS
 Defintion, 61
 PMIx_Proc_state_string, 9
 Defintion, 84
 pmix_proc_state_t, 24, 51, 84
 Defintion, 24
 PMIX_PROC_STATE_TERM_NON_ZERO,
 25
 PMIX_PROC_STATE_TERM_WO_SYNC,
 25
 PMIX_PROC_STATE_TERMINATE, 25
 PMIX_PROC_STATE_TERMINATED, 25
 PMIX_PROC_STATE_UNDEF, 25
 PMIX_PROC_STATE_UNTERMINATED,
 25
 pmix_proc_t, 22–24, 39, 40, 51, 57, 63, 79,
 83, 89, 91, 93, 103, 113–115, 129,
 171, 174, 175, 180, 181, 192,
 202–205, 211–213, 215, 217, 219,
 221, 223, 225, 229, 232, 237, 240,
 243, 245, 248, 250
 Defintion, 22
 PMIX_PROC_TERMINATED, 19

PMIX_PROC_URI, 154
 Defintion, **57**
 PMIX_PROCDIR
 Defintion, **56**
 PMIx_Process_monitor_nb, 9, 71, 148, 164
 Defintion, **161**
 PMIX_PROCID, 152, 153, 192
 Defintion, **56**
 PMIX_PROGRAMMING_MODEL
 Defintion, **54**
 PMIx_Publish, 8, 28, 61, 118, 119, 121, 220
 Defintion, **117**
 PMIx_Publish_nb, 8, 121
 Defintion, **119**
 PMIx_Put, 8, 27–29, 82, 100, 103, 106, 112, 114, 139, 155, 205, 218
 Defintion, **99**
 PMIX_QUERY, 51
 PMIX_QUERY_ALLOC_STATUS, 154, 241
 Defintion, **67**
 PMIX_QUERY_AUTHORIZATIONS
 Defintion, **67**
 PMIX_QUERY_CONSTRUCT
 Defintion, **43**
 PMIX_QUERY_CREATE
 Defintion, **44**
 PMIX_QUERY_DEBUG_SUPPORT, 153, 241
 Defintion, **67**
 PMIX_QUERY_DESTRUCT
 Defintion, **43**
 PMIX_QUERY_FREE
 Defintion, **44**
 PMIx_Query_info_nb, 9, 10, 43, 66, 111, 139, 148, 154, 155
 Defintion, **150**
 PMIX_QUERY_JOB_STATUS, 153, 240
 Defintion, **67**
 PMIX_QUERY_LOCAL_ONLY, 241
 Defintion, **67**
 PMIX_QUERY_LOCAL_PROC_TABLE, 153, 241
 Defintion, **67**
 PMIX_QUERY_MEMORY_USAGE, 153, 241
 Defintion, **67**
 PMIX_QUERY_NAMESPACES, 153, 240
 Defintion, **67**
 PMIX_QUERY_PARTIAL_SUCCESS, 19
 PMIX_QUERY_PROC_TABLE, 153, 241
 Defintion, **67**
 PMIX_QUERY_QUEUE_LIST, 153, 240
 Defintion, **67**
 PMIX_QUERY_QUEUE_STATUS, 153, 240
 Defintion, **67**
 PMIX_QUERY_REFRESH_CACHE, 151, 154, 155
 Defintion, **66**
 PMIX_QUERY_REPORT_AVG, 153, 241
 Defintion, **67**
 PMIX_QUERY_REPORT_MINMAX, 153, 241
 Defintion, **67**
 PMIX_QUERY_SPAWN_SUPPORT, 153, 241
 Defintion, **67**
 pmix_query_t, 43, 44, 152, 153, 155, 240, 241
 Defintion, **43**
 PMIX_RANGE, 118, 121, 122, 124, 126, 128, 171, 219, 222, 224, 238
 Defintion, **61**
 PMIX_RANGE_CUSTOM, 28
 PMIX_RANGE_GLOBAL, 28
 PMIX_RANGE_LOCAL, 28
 PMIX_RANGE_NAMESPACE, 28
 PMIX_RANGE_PROC_LOCAL, 28
 PMIX_RANGE_RM, 28
 PMIX_RANGE_SESSION, 28
 PMIX_RANGE_UNDEF, 28
 PMIX_RANK, 152, 153, 190
 Defintion, **56**
 PMIX_RANK_LOCAL_NODE, 22
 pmix_rank_t, 22, 24, 51

Defintion, [22](#)
 PMIX_RANK_UNDEF, [22](#)
 PMIX_RANK_WILDCARD, [22](#)
 PMIX_RANKBY, [132](#), [137](#), [191](#), [227](#)
 Defintion, [65](#)
 PMIx_Register_event_handler, [9](#), [77](#), [148](#)
 Defintion, [169](#)
 PMIX_REGISTER_NODATA, [189](#)
 Defintion, [53](#), [61](#)
 pmix_release_cbfunc_t, [72](#)
 Defintion, [72](#)
 PMIX_REMOTE, [27](#)
 PMIX_REPORT_BINDINGS, [134](#), [138](#), [228](#)
 Defintion, [66](#)
 PMIX_REQUESTOR_IS_CLIENT, [131](#),
 [135](#)
 Defintion, [54](#)
 PMIX_REQUESTOR_IS_TOOL, [131](#), [136](#)
 Defintion, [54](#)
 PMIx_Resolve_nodes, [8](#)
 Defintion, [149](#)
 PMIx_Resolve_peers, [8](#)
 Defintion, [149](#)
 PMIX_RM_NAME
 Defintion, [68](#)
 PMIX_RM_VERSION
 Defintion, [69](#)
 PMIX_SCOPE, [51](#)
 PMIx_Scope_string, [9](#)
 Defintion, [84](#)
 pmix_scope_t, [27](#), [51](#), [84](#), [100](#)
 Defintion, [27](#)
 PMIX_SCOPE_UNDEF, [27](#)
 PMIX_SEND_HEARTBEAT
 Defintion, [71](#)
 pmix_server_abort_fn_t
 Defintion, [213](#)
 pmix_server_alloc_fn_t
 Defintion, [245](#)
 pmix_server_client_connected_fn_t, [74](#), [202](#),
 [203](#), [211](#)
 Defintion, [210](#)
 pmix_server_client_finalized_fn_t, [212](#), [213](#)
 Defintion, [212](#)
 pmix_server_connect_fn_t, [231](#), [233](#)
 Defintion, [229](#)
 PMIx_server_deregister_client, [8](#)
 Defintion, [203](#)
 pmix_server_deregister_events_fn_t
 Defintion, [235](#)
 PMIx_server_deregister_nspace, [8](#), [203](#)
 Defintion, [200](#)
 pmix_server_disconnect_fn_t, [233](#)
 Defintion, [231](#)
 pmix_server_dmodex_req_fn_t, [72](#)
 Defintion, [217](#)
 PMIx_server_dmodex_request, [9](#), [81](#), [82](#),
 [205](#)
 Defintion, [204](#)
 PMIX_SERVER_ENABLE_MONITORING
 Defintion, [53](#)
 pmix_server_fencenb_fn_t, [72](#), [216](#)
 Defintion, [214](#)
 PMIx_server_finalize, [8](#)
 Defintion, [97](#)
 PMIX_SERVER_HOSTNAME
 Defintion, [53](#)
 PMIx_server_init, [8](#), [87](#), [209](#)
 Defintion, [95](#)
 pmix_server_job_control_fn_t
 Defintion, [247](#)
 pmix_server_listener_fn_t
 Defintion, [238](#)
 pmix_server_log_fn_t
 Defintion, [243](#)
 pmix_server_lookup_fn_t
 Defintion, [220](#)
 pmix_server_module_t, [95](#), [97](#), [209](#)
 Defintion, [209](#)
 pmix_server_monitor_fn_t
 Defintion, [250](#)
 pmix_server_notify_event_fn_t, [80](#)
 Defintion, [237](#)
 PMIX_SERVER_NAMESPACE, [95](#), [191](#)
 Defintion, [53](#)
 PMIX_SERVER_PIDINFO, [92](#), [93](#)

Defintion, [53](#)
 pmix_server_publish_fn_t
 Defintion, [218](#)
 pmix_server_query_fn_t
 Defintion, [239](#)
 PMIX_SERVER_RANK, [95](#), [191](#)
 Defintion, [53](#)
 PMIx_server_register_client, [8](#), [203](#), [211](#),
 [212](#)
 Defintion, [201](#)
 pmix_server_register_events_fn_t
 Defintion, [233](#)
 PMIx_server_register_nspace, [8](#), [10](#), [14](#), [58](#),
 [74](#), [193](#), [196](#)
 Defintion, [188](#)
 PMIX_SERVER_REMOTE_CONNECTIONS,
 [97](#)
 Defintion, [52](#)
 PMIx_server_setup_application, [9](#), [80](#), [81](#),
 [208](#)
 Defintion, [206](#)
 PMIx_server_setup_fork, [9](#)
 Defintion, [204](#)
 PMIx_server_setup_local_support, [9](#)
 Defintion, [207](#)
 pmix_server_spawn_fn_t, [73](#)
 Defintion, [224](#)
 PMIX_SERVER_SYSTEM_SUPPORT, [96](#)
 Defintion, [52](#)
 PMIX_SERVER_TMPDIR, [95](#)
 Defintion, [52](#)
 pmix_server_tool_connection_fn_t
 Defintion, [242](#)
 PMIX_SERVER_TOOL_SUPPORT, [95](#)
 Defintion, [52](#)
 pmix_server_unpublish_fn_t
 Defintion, [222](#)
 PMIX_SERVER_URI, [91](#), [93](#), [154](#)
 Defintion, [53](#)
 PMIX_SESSION_ID, [57](#), [58](#), [101](#), [105](#), [108](#),
 [151](#), [191](#), [196](#)
 Defintion, [57](#)
 PMIX_SESSION_INFO, [101](#), [105](#), [107](#), [151](#)
 Defintion, [57](#)
 PMIX_SESSION_INFO_ARRAY, [59](#), [189](#),
 [196](#)
 Defintion, [58](#)
 PMIX_SET_ENVAR
 Defintion, [69](#)
 PMIX_SET_SESSION_CWD, [132](#), [136](#), [226](#)
 Defintion, [65](#)
 pmix_setup_application_cbfunc_t, [206](#)
 Defintion, [80](#)
 PMIX_SINGLE_LISTENER, [88](#)
 Defintion, [54](#)
 PMIX_SIZE, [50](#)
 PMIX_SOCKET_MODE, [88](#), [92](#), [96](#)
 Defintion, [54](#)
 PMIx_Spawn, [8](#), [41](#), [55](#), [64](#), [68](#), [69](#), [130](#), [131](#),
 [135](#), [136](#), [138](#), [193](#), [204](#), [224](#), [229](#)
 Defintion, [130](#)
 pmix_spawn_cbfunc_t, [73](#), [135](#), [225](#)
 Defintion, [73](#)
 PMIx_Spawn_nb, [8](#), [41](#), [73](#)
 Defintion, [135](#)
 PMIX_SPAWNED, [131](#), [135](#), [136](#), [226](#)
 Defintion, [55](#)
 PMIX_STATUS, [51](#)
 pmix_status_t, [17](#), [32](#), [51](#), [76](#), [77](#), [79](#), [81–84](#),
 [170](#), [174](#), [234](#), [236](#), [237](#)
 Defintion, [17](#)
 PMIX_STDIN_TGT, [133](#), [137](#), [227](#)
 Defintion, [65](#)
 PMIx_Store_internal, [9](#)
 Defintion, [106](#)
 PMIX_STRING, [50](#)
 PMIX_SUCCESS, [18](#)
 PMIX_SYSTEM_TMPDIR, [95](#)
 Defintion, [52](#)
 PMIX_TAG_OUTPUT, [133](#), [137](#), [227](#)
 Defintion, [66](#)
 PMIX_TCP_DISABLE_IPV4, [89](#), [93](#), [96](#)
 Defintion, [55](#)
 PMIX_TCP_DISABLE_IPV6, [89](#), [93](#), [96](#)
 Defintion, [55](#)
 PMIX_TCP_IF_EXCLUDE, [88](#), [92](#), [96](#)

Definition, [55](#)
 PMIX_TCP_IF_INCLUDE, [88](#), [92](#), [96](#)
 Definition, [55](#)
 PMIX_TCP_IPV4_PORT, [89](#), [92](#), [96](#)
 Definition, [55](#)
 PMIX_TCP_IPV6_PORT, [89](#), [92](#), [96](#)
 Definition, [55](#)
 PMIX_TCP_REPORT_URI, [88](#), [92](#), [96](#)
 Definition, [54](#)
 PMIX_TCP_URI, [92](#), [93](#)
 Definition, [54](#)
 PMIX_TDIR_RMCLEAN
 Definition, [56](#)
 PMIX_THREADING_MODEL
 Definition, [54](#)
 PMIX_TIME, [51](#)
 PMIX_TIME_REMAINING, [148](#), [154](#), [241](#)
 Definition, [67](#)
 PMIX_TIMEOUT, [3](#), [11](#), [102](#), [103](#), [105](#), [106](#),
 [113](#), [114](#), [116](#), [118](#), [121–126](#), [128](#),
 [140](#), [141](#), [143](#), [145](#), [147](#), [216](#), [218](#),
 [220](#), [222](#), [224](#), [228](#), [230](#), [233](#)
 Definition, [60](#)
 PMIX_TIMESTAMP_OUTPUT, [133](#), [137](#),
 [227](#)
 Definition, [66](#)
 PMIX_TIMEVAL, [51](#)
 PMIX_TMPDIR, [56](#)
 Definition, [56](#)
 pmix_tool_connection_cbfunc_t, [242](#)
 Definition, [83](#)
 PMIX_TOOL_DO_NOT_CONNECT, [91](#),
 [93](#)
 Definition, [53](#)
 PMIx_tool_finalize, [9](#)
 Definition, [94](#)
 PMIx_tool_init, [9](#), [53](#), [87](#), [94](#)
 Definition, [91](#)
 PMIX_TOOL_NAMESPACE, [91](#)
 Definition, [53](#)
 PMIX_TOOL_RANK, [91](#)
 Definition, [53](#)
 PMIX_TOPOLOGY
 Definition, [60](#)
 PMIX_TOPOLOGY_SIGNATURE
 Definition, [60](#)
 PMIX_UINT, [50](#)
 PMIX_UINT16, [51](#)
 PMIX_UINT32, [51](#)
 PMIX_UINT64, [51](#)
 PMIX_UINT8, [50](#)
 PMIX_UNDEF, [50](#)
 PMIX_UNIV_SIZE, [10](#), [103](#), [106](#), [107](#), [189](#),
 [196](#)
 Definition, [58](#)
 PMIx_Unpublish, [8](#), [127](#), [128](#)
 Definition, [125](#)
 PMIx_Unpublish_nb, [8](#)
 Definition, [127](#)
 PMIX_UNSET_ENVAR
 Definition, [69](#)
 PMIX_USERID, [118](#), [120](#), [122](#), [124](#), [126](#),
 [128](#), [152](#), [157](#), [159](#), [162](#), [165](#),
 [219–225](#), [234](#), [240](#), [242](#), [244](#), [246](#),
 [248](#), [251](#)
 Definition, [53](#)
 PMIX_USOCK_DISABLE, [88](#), [96](#)
 Definition, [54](#)
 PMIX_VALUE, [51](#)
 pmix_value_cbfunc_t, [75](#)
 Definition, [75](#)
 PMIX_VALUE_CONSTRUCT
 Definition, [30](#)
 PMIX_VALUE_CREATE
 Definition, [30](#)
 PMIX_VALUE_DESTRUCT
 Definition, [30](#)
 PMIX_VALUE_FREE
 Definition, [31](#)
 PMIX_VALUE_LOAD
 Definition, [31](#)
 pmix_value_t, [29–32](#), [51](#), [75](#), [99](#), [100](#)
 Definition, [29](#)
 PMIX_VALUE_XFER
 Definition, [32](#)
 PMIX_VERSION_INFO

Definition, [54](#)
PMIX_WAIT, [122–124](#), [222](#)
Definition, [60](#)
PMIX_WDIR, [131](#), [136](#), [226](#)
Definition, [64](#)
rank, [109](#), [197](#)
Definition, [12](#)
resource manager
Definition, [13](#)

session, [10](#), [57](#), [58](#), [99](#), [107](#), [155](#), [193](#)
Definition, [12](#)
slot
Definition, [12](#)
slots
Definition, [12](#)
workflow
Definition, [12](#)