PMIx Standard
Administrative Steering Committee (ASC)
Q4 2019 Meeting

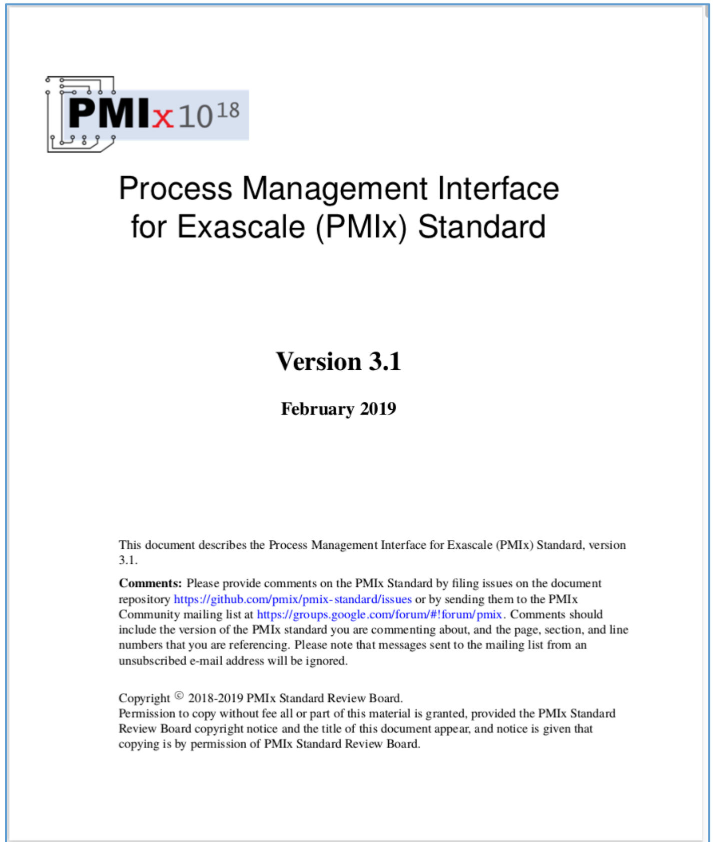

Process Management Interface
for Exascale (PMIx) Standard

Version 3.1

February 2019

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 3.1.

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.

Joshua Hursey (IBM)

Co-Chair 2019-2021

https://github.com/pmix/pmix-standard
https://github.com/pmix/governance

Agenda

https://github.com/pmix/pmix-standard/wiki/ASC-Q4-2019-Meeting

What is PMIx?

PMIx is a standard API providing libraries and programming models with portable and well-defined access to commonly available system services.

- PMIx works as a messenger between these pieces of software, not a doer.
  - Facilitates the interaction between applications, tools, runtime environments.
- Open, community driven standard.
- Use cases: (summarized list)
  - **Process wire-up** via either business card exchange or “instant on” (where supported)
  - **Tool connections** including debugger support
  - **Event notification** used by fault tolerant libraries
  - Application/Job/Node **environment discovery**
  - Job scheduler interaction

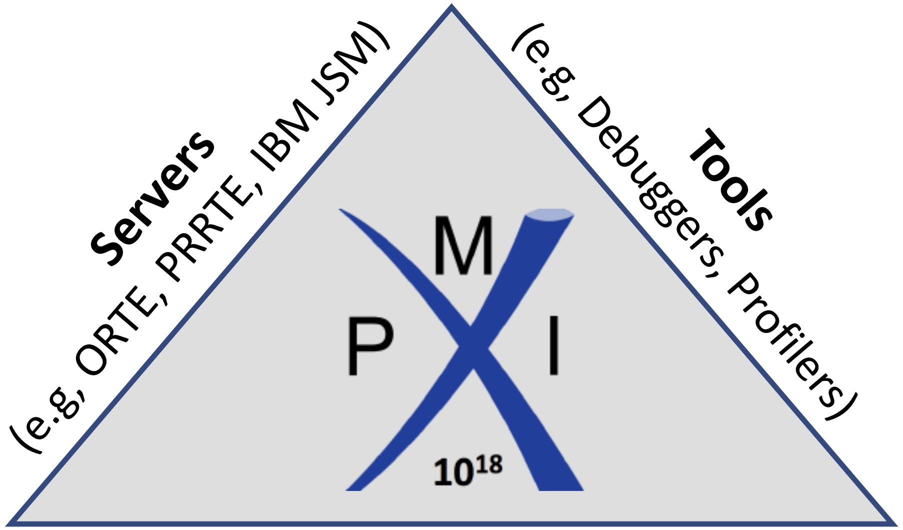

Clients
(e.g, MPI, OpenSHMEM)

https://github.com/pmix/pmix-standard
https://github.com/pmix/governance

PMIx 10^18

What is PMIx?

- PMIx enabled server does not need to implement 100% of the PMIx interface.
  - Implementation level can be driven by the range of use cases they need to support.

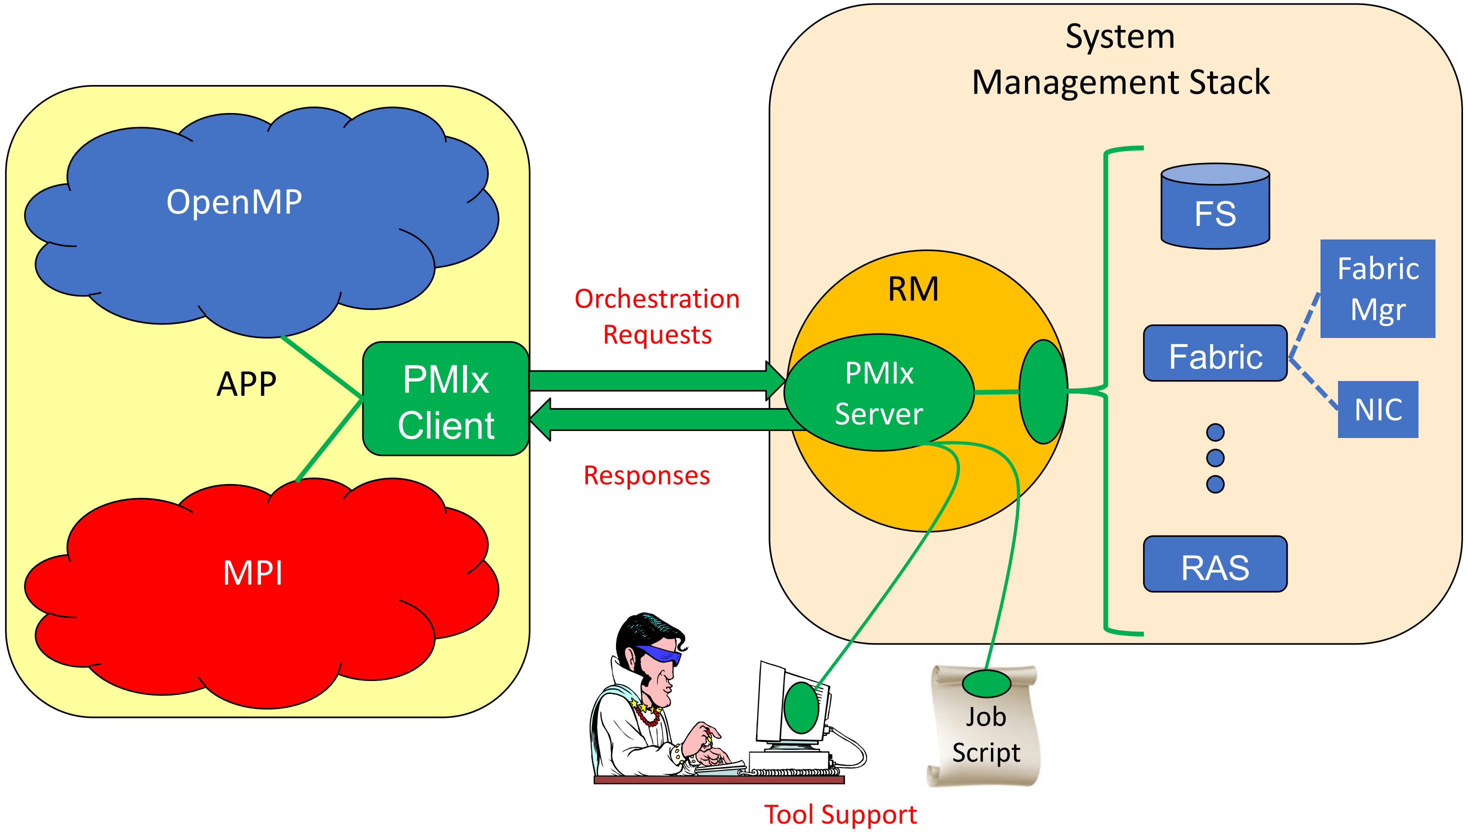

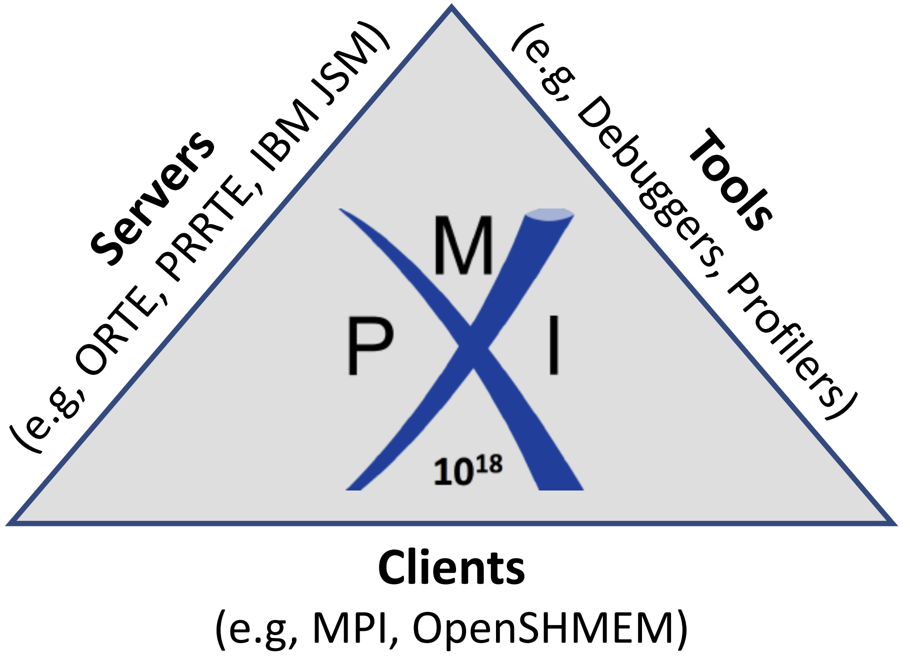

PMIx 10^18

https://github.com/pmix/pmix-standard
https://github.com/pmix/governance

A Brief History

Artem Polyakov, et.al., “A Performance Analysis and Optimization of PMIx-based HPC Software Stacks”, EuroMPI 2019.

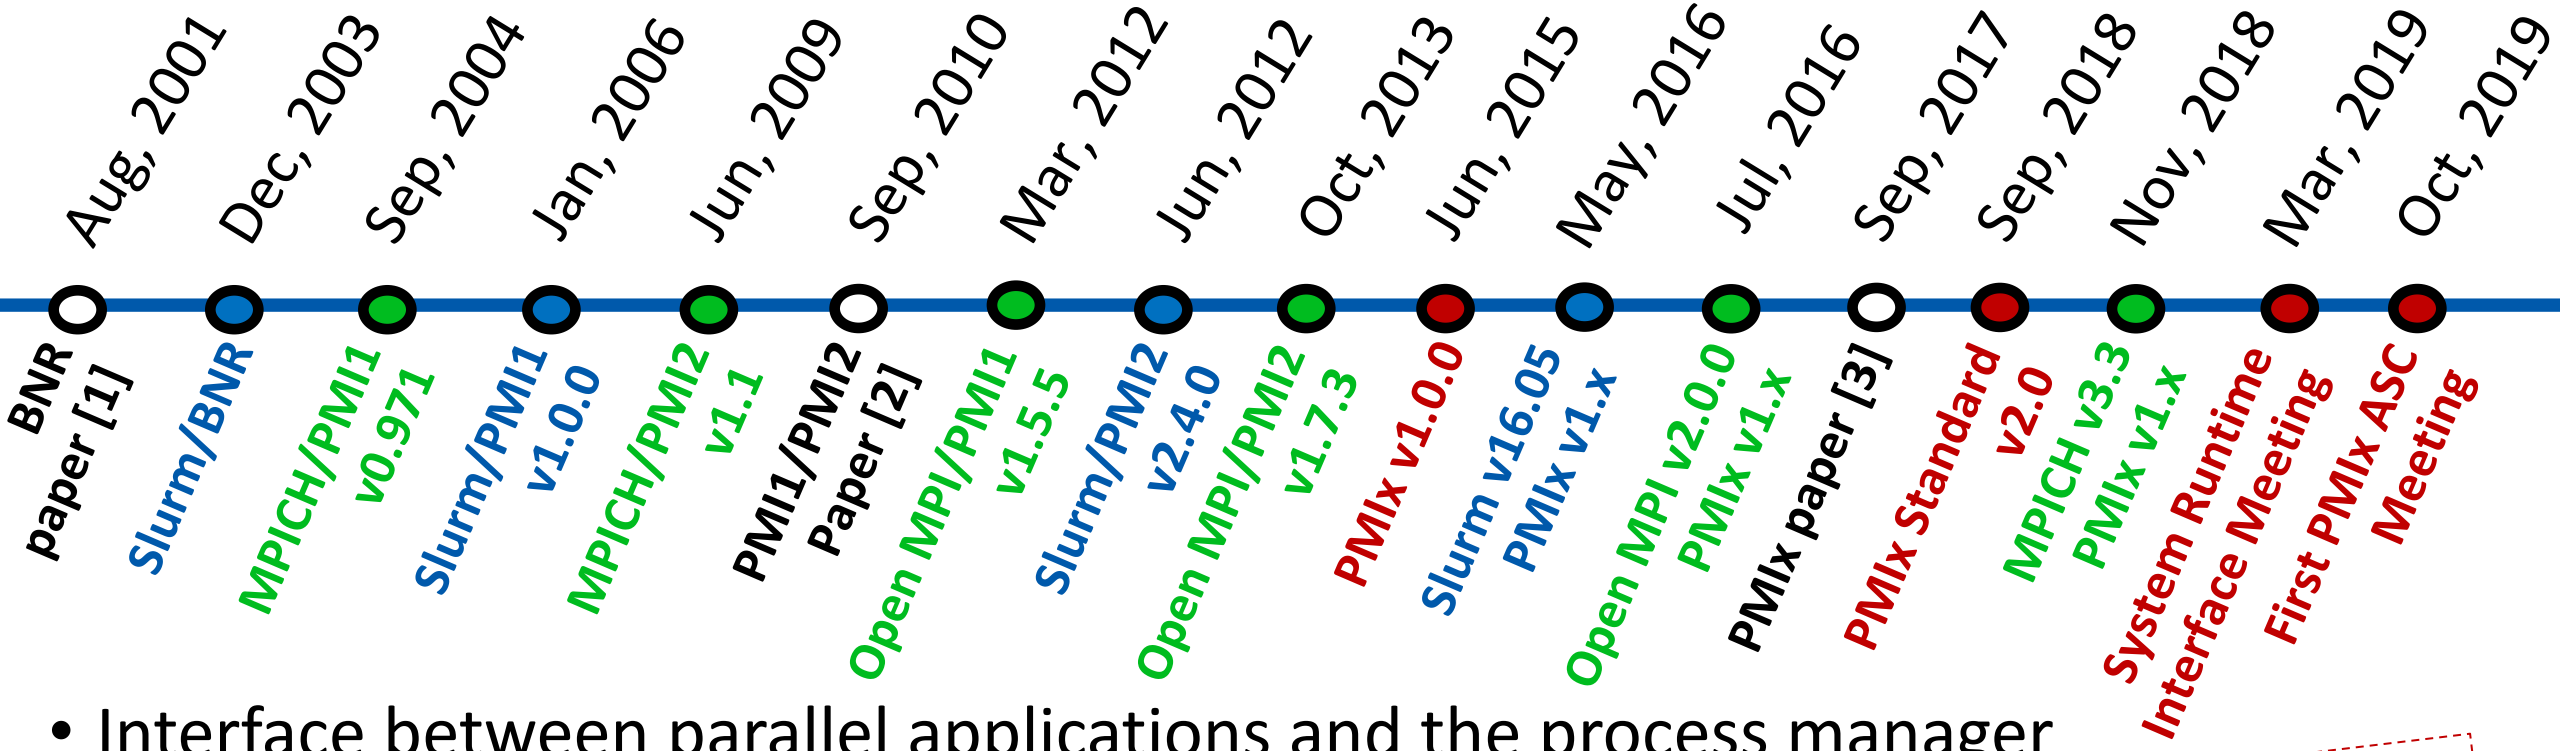

Aug, 2001 — BNR paper [1]
Dec, 2003 — Slurm/BNR
Sep, 2004 — MPICH/PMI1 v0.971
Jan, 2006 — Slurm/PMI1 v1.0.0
Jun, 2009 — MPICH/PMI2 v1.1
Sep, 2010 — PMI1/PMI2 Paper [2]
Mar, 2012 — Open MPI/PMI1 v1.5.5
Jun, 2012 — Slurm/PMI2 v2.4.0
Oct, 2013 — Open MPI/PMI2 v1.7.3
Jun, 2015 — PMIx v1.0.0
May, 2016 — Slurm v16.05 PMIx v1.x
Jul, 2016 — Open MPI v2.0.0 PMIx v1.x
Sep, 2017 — PMIx paper [3]
Sep, 2018 — PMIx Standard v2.0
Nov, 2018 — MPICH v3.3 PMIx v1.x
Mar, 2019 — System Runtime Interface Meeting
Oct, 2019 — First PMIx ASC Meeting

- Interface between parallel applications and the process manager continues to be a critical part of the HPC software ecosystem.
- The community needs to balance active innovation with interface stability.

Put
Commit
Fence
Get

[1] B. Toonen, et al. Interfacing parallel jobs to process managers, 10th IEEE HPDC, 2001, San Francisco, USA
[2] Pavan Balaji, et al. PMI: A Scalable Parallel Process management Interface for Extreme-scale Systems. EuroMPI’10, Stuttgart, Germany.
[3] R.H. Castain, et al. PMIx: Process Management for Exascale Environments, EuroMPI 2017, Chicago, USA

PMIx Example Client Program
OpenPMIx 3.1.4, Open MPI 4.0.1

```c
#include <stdio.h>
#include <pmix.h>
int main(int argc, char **argv) {
    pmix_proc_t myproc, proc_wildcard;
    pmix_value_t value;
    pmix_value_t *val = &value;
    int job_size;
    char hostname[256];

    PMIx_Init(&myproc, NULL, 0);
    PMIX_PROC_CONSTRUCT(&proc_wildcard);
    (void)strncpy(proc_wildcard.nspace, myproc.nspace, PMIX_MAX_NSLEN);
    proc_wildcard.rank = PMIX_RANK_WILDCARD;

    PMIx_Get(&proc_wildcard, PMIX_JOB_SIZE, NULL, 0, &val);
    job_size = val->data.uint32;
    PMIx_Get(&myproc, PMIX_HOSTNAME, NULL, 0, &val);
    strncpy(hostname, val->data.string, 256);

    printf("%d/%d) Hello World from %s\n", myproc.rank, job_size, hostname);

    PMIx_Get(&proc_wildcard, PMIX_LOCALLDR, NULL, 0, &val); // Lowest rank on this node
    printf("%d/%d) Lowest Local Rank: %d\n", myproc.rank, job_size, val->data.rank);
    PMIx_Get(&proc_wildcard, PMIX_LOCAL_SIZE, NULL, 0, &val); // Number of ranks on this node
    printf("%d/%d) Local Ranks: %d\n", myproc.rank, job_size, val->data.uint32);

    PMIx_Fence(&proc_wildcard, 1, NULL, 0); // Synchronize processes (not required, just for demo)
    PMIx_Finalize(NULL, 0); // Cleanup
    return 0;
}
```

```
$$ gcc -I$PMIX_ROOT/include/ -L$PMIX_ROOT/lib –lpmix test.c
$$ mpirun -npernode 2 --host f3n17:2,f3n19:2 ./a.out
0/4) Hello World from f3n17
0/4) Lowest Local Rank: 0
0/4) Local Ranks: 2
1/4) Hello World from f3n17
1/4) Lowest Local Rank: 0
1/4) Local Ranks: 2
2/4) Hello World from f3n19
2/4) Lowest Local Rank: 2
2/4) Local Ranks: 2
3/4) Hello World from f3n19
3/4) Lowest Local Rank: 2
3/4) Local Ranks: 2
```

PMIx Administrative Steering Committee (ASC)

The open governance body for the PMIx Standard with broad representation.

- Mission:
  - Define the mission of the PMIx Standard
  - Plan release timelines for versions of the PMIx Standard
  - Assemble working groups for areas of interest
  - Vote on all issues (1 vote per Member organization)
  - Vote on ASC leadership positions
- Meetings:
  - Quarterly ASC meetings to vote on items that move the PMIx Standard forward.
    - One quarter per year may be face-to-face.
  - Regular PMIx Standard teleconference to drive progress between quarterly meetings.
  - Working Group meetings to drive specific areas of interest and need.

PMIx ASC Membership and Leadership Positions

- Participation:
  - **Participant**: Informal, non-voting individual
  - **Member**: Formal, voting organization or individual
    - Must attend 2 quarterly meetings in last 12 months
    - Membership status must be requested (or nominated) and voted on in a quarterly meeting.
    - Only one vote per Member organization. Must designate a Representative with voting authority.

Member attendance requirement waved for this first meeting.

- Leadership:
  - Elections during last quarterly meeting of the year. Start term at beginning of next year.
  - **Co-chairs**:
    - Ralph Castain, Intel (2018-2020) (even year)
    - Joshua Hursey, IBM (2019-2021) (odd year)
  - **Secretaries** (odd and even year)
  - **Release Managers** (2 per major release)

PMIx Standard Voting and Release Timeline

- Voting:
  - Goal is to reach consensus on all, if not most, issues.
  - Each ASC Member organization receives 1 vote.
  - 2/3 majority required to pass. A 50% quorum must be present.
  - Straw Poll: Informal voting mechanism to gauge general support for an issue.
- PMIx Standard release timeline structure
  - Minor releases (`A.x`) can occur as quickly as 1 per quarter.
    - Release manager proposes a minor release 1 month before next quarterly meeting for approval.
  - Major releases (`X.0`) can occur as quickly as 1 per year at face-to-face meeting.
    - Release manager proposes a major release 2 quarters before the face-to-face quarterly meeting.
    - If pushed back then can be approved in any of the next quarterly meetings.

Modifying the PMIx ASC Governance Document

Vote in Two Consecutive Quarterly Meetings

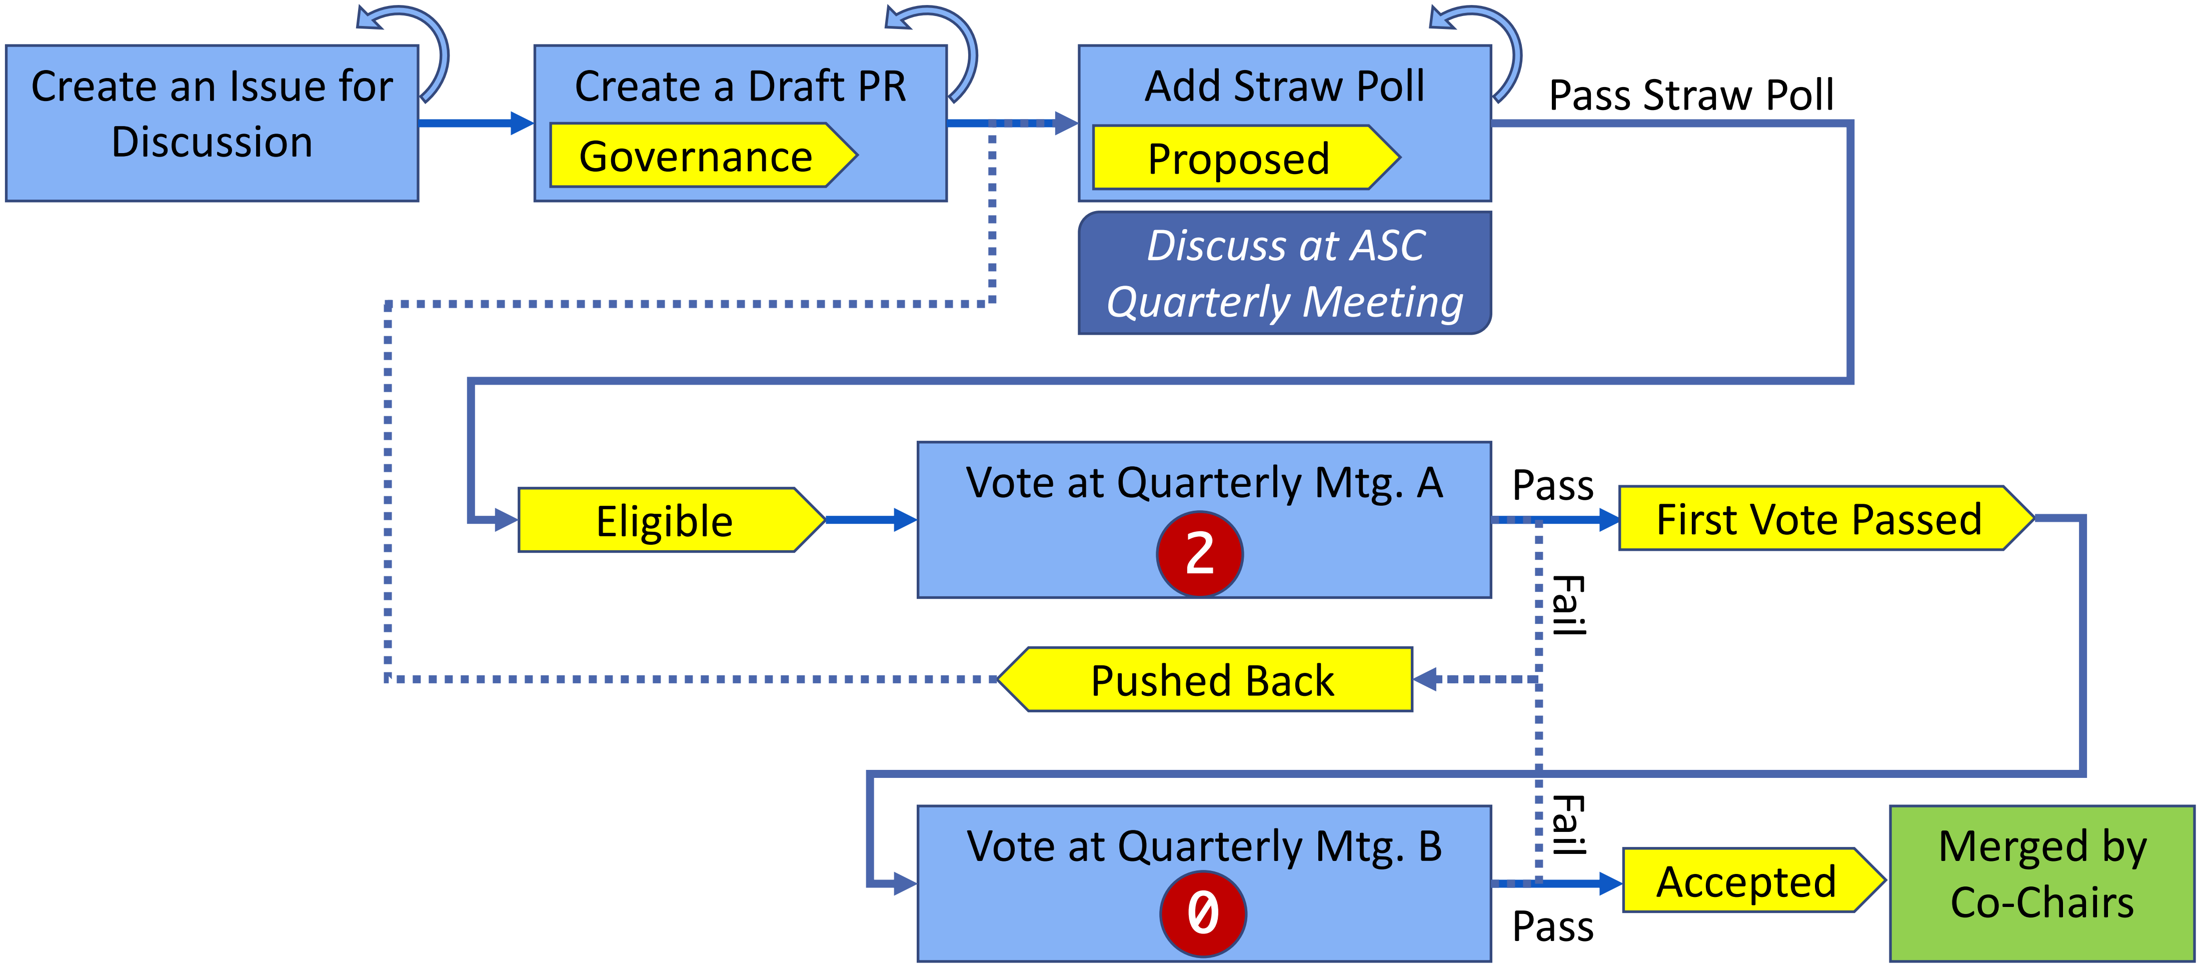

Create an Issue for Discussion

Create a Draft PR

Governance

Add Straw Poll

Proposed

Discuss at ASC Quarterly Meeting

Pass Straw Poll

Eligible

Vote at Quarterly Mtg. A

2

Pass

First Vote Passed

Fail

Pushed Back

Vote at Quarterly Mtg. B

0

Fail

Accepted

Pass

Merged by Co-Chairs

PMIx 10^{18}

PMIx Standard Process

- Rely on GitHub Issues and Pull Requests
  - Archives conversation around a change – for future us.
  - Facilitates asynchronous conversations from geographically distributed participants.
  - Consensus building medium for resolving questions and dissenting opinions.
- PMIx APIs and attributes are grouped into three classes:
  - **Provisional**: Included in the standard, but subject to change. *Everything new starts here*.
    - Added in any minor release
    - Modified after 2 minor releases (1st minor release = warning, 2nd minor release = modification)
  - **Stable**: Permanent, non-changing parts. Subject only to deprecation.
    - Elevate from Provisional to Stable at a major release
  - **Deprecated**: Items slated for removal from the standard.
    - Provisional items may be deprecated and removed after 2 minor releases
    - Stable items may be deprecated and removed after 2 major releases

PMIx Standard: Provisional Items

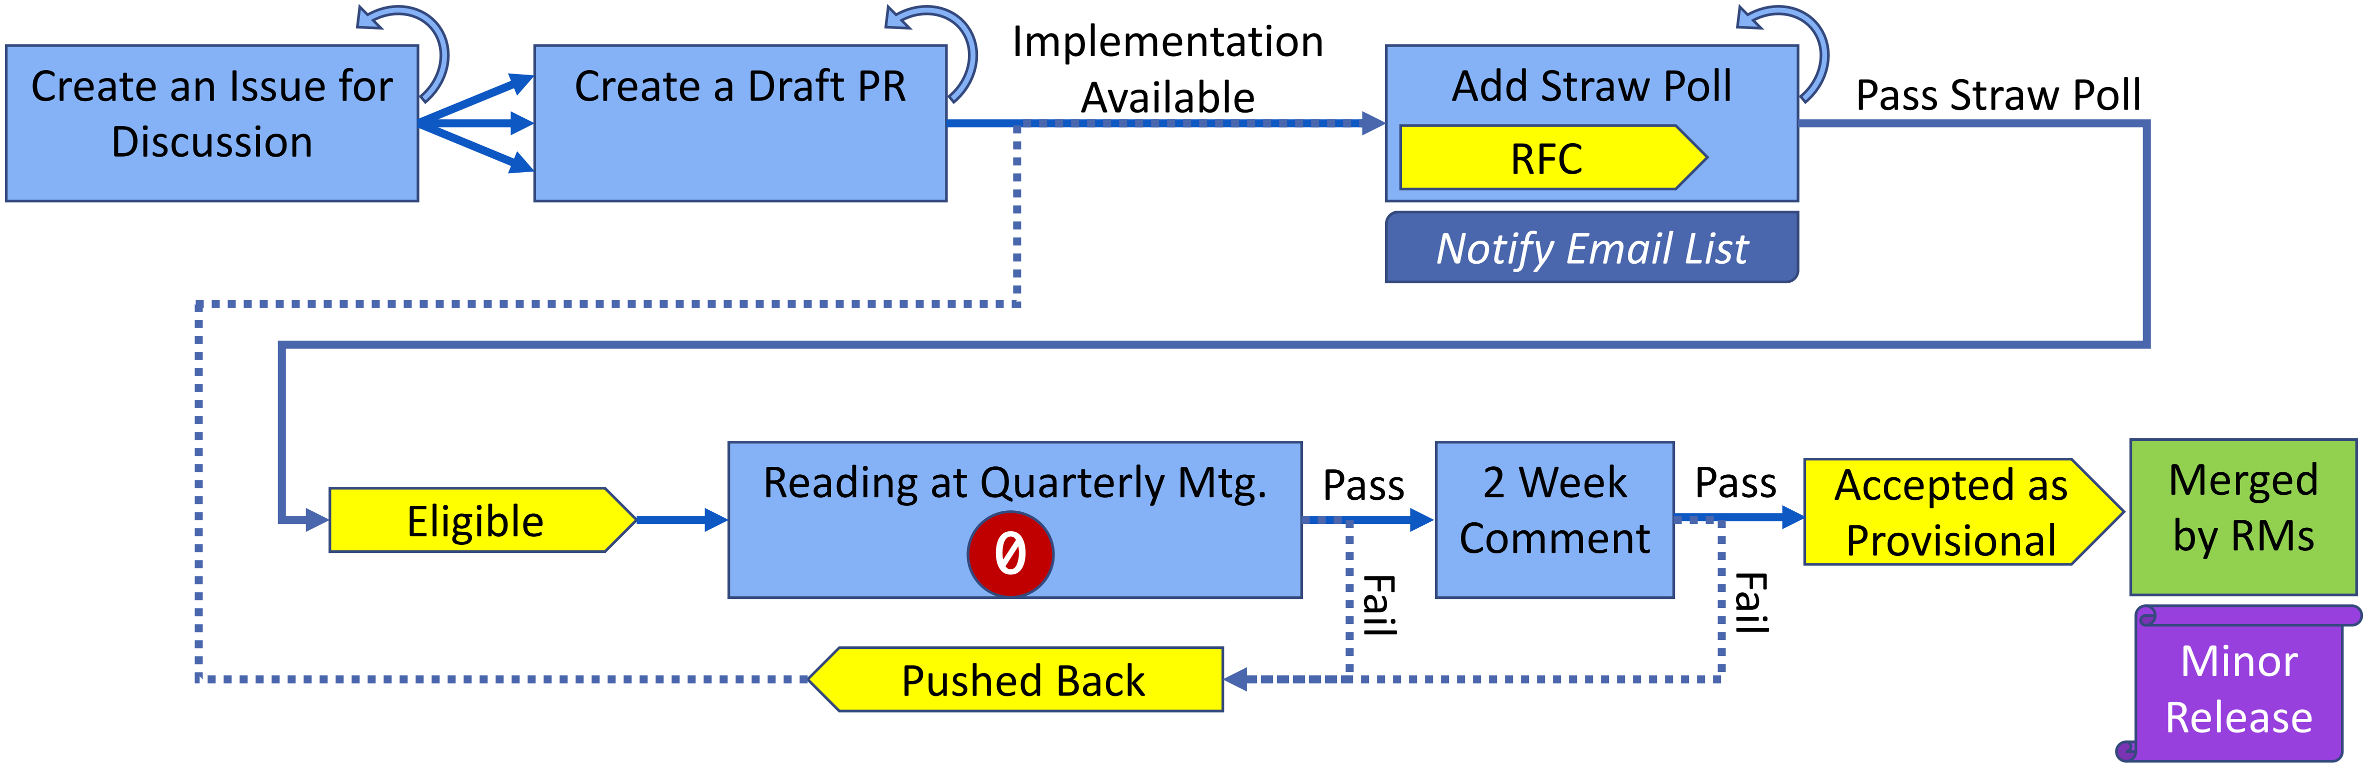

PMIx Standard: Elevating Provisional to Stable

Vote in Two Consecutive Quarterly Meetings

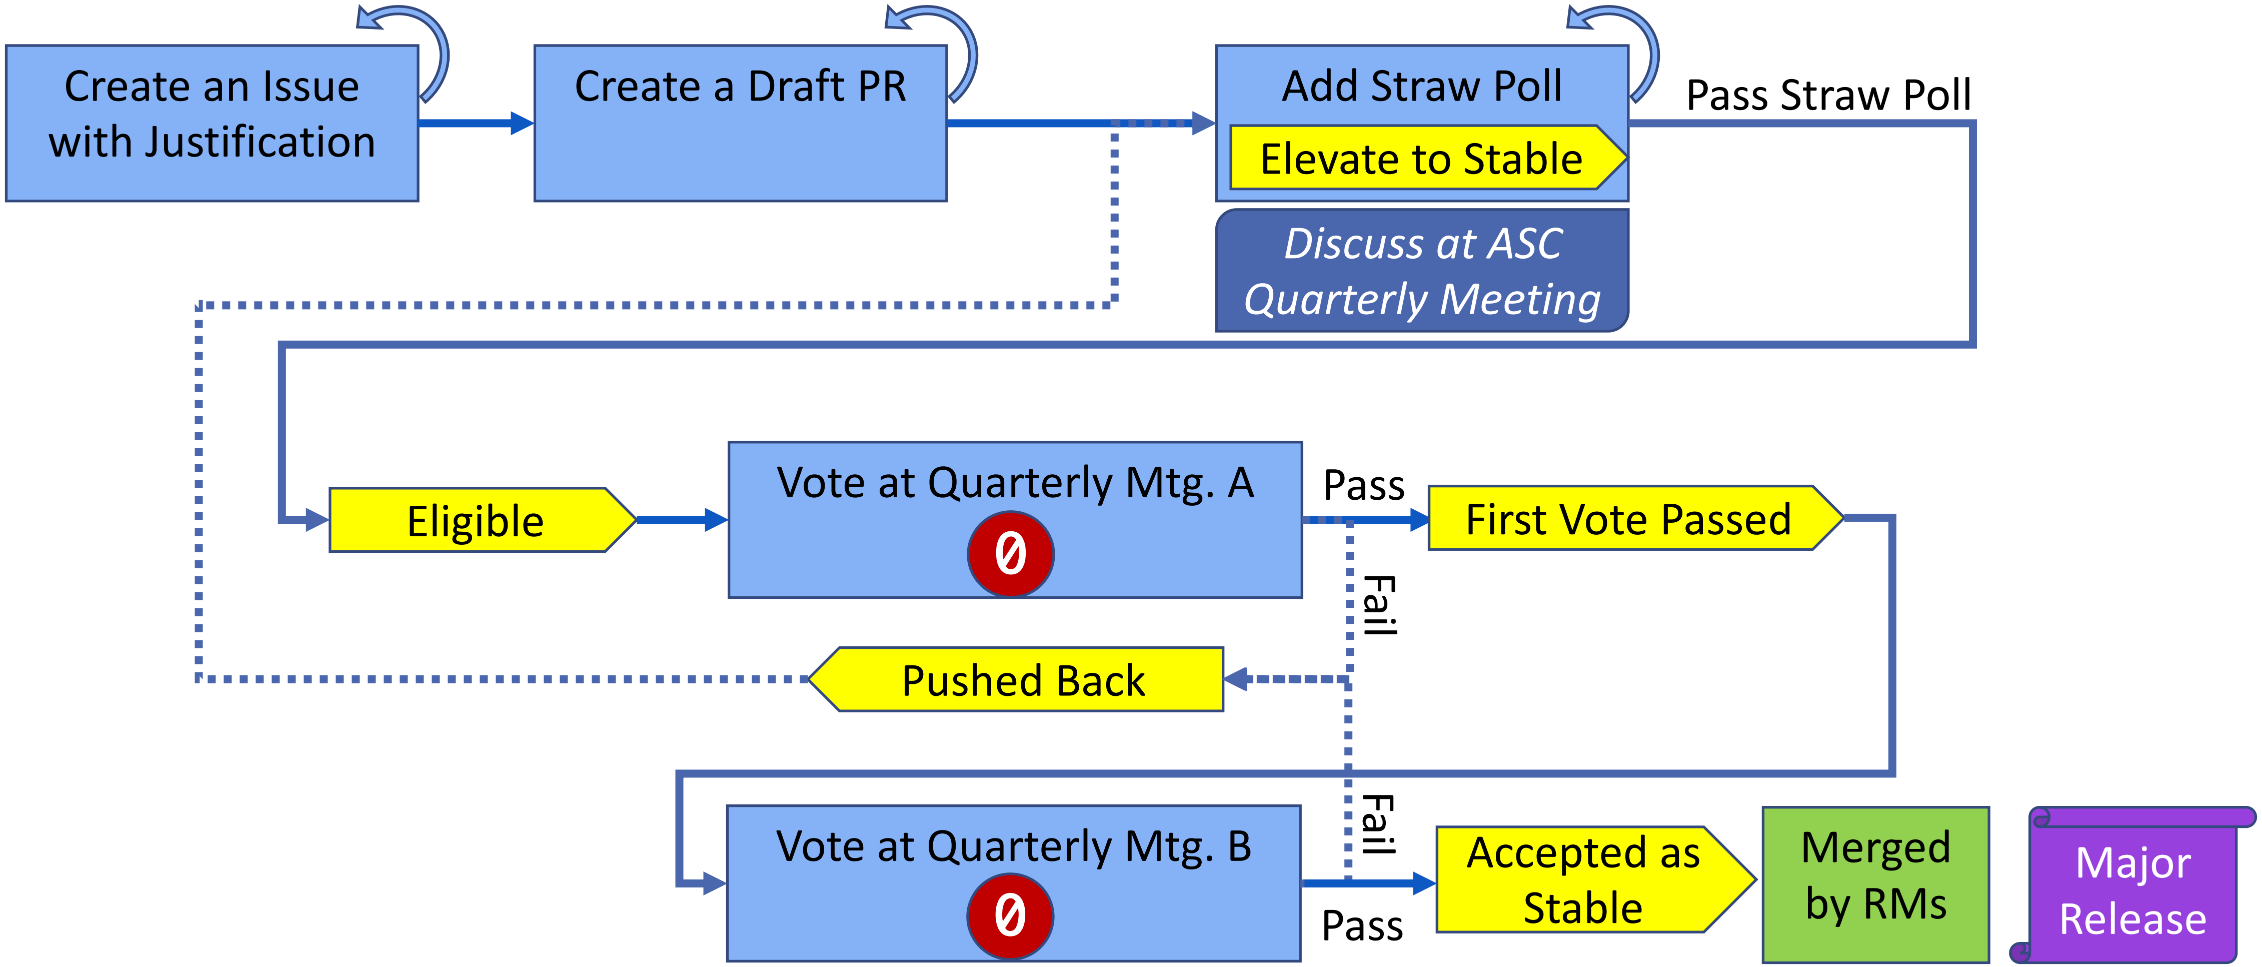

PMIx 10^{18}

PMIx Standard: Deprecation from Stable

Vote in Two Consecutive Quarterly Meetings

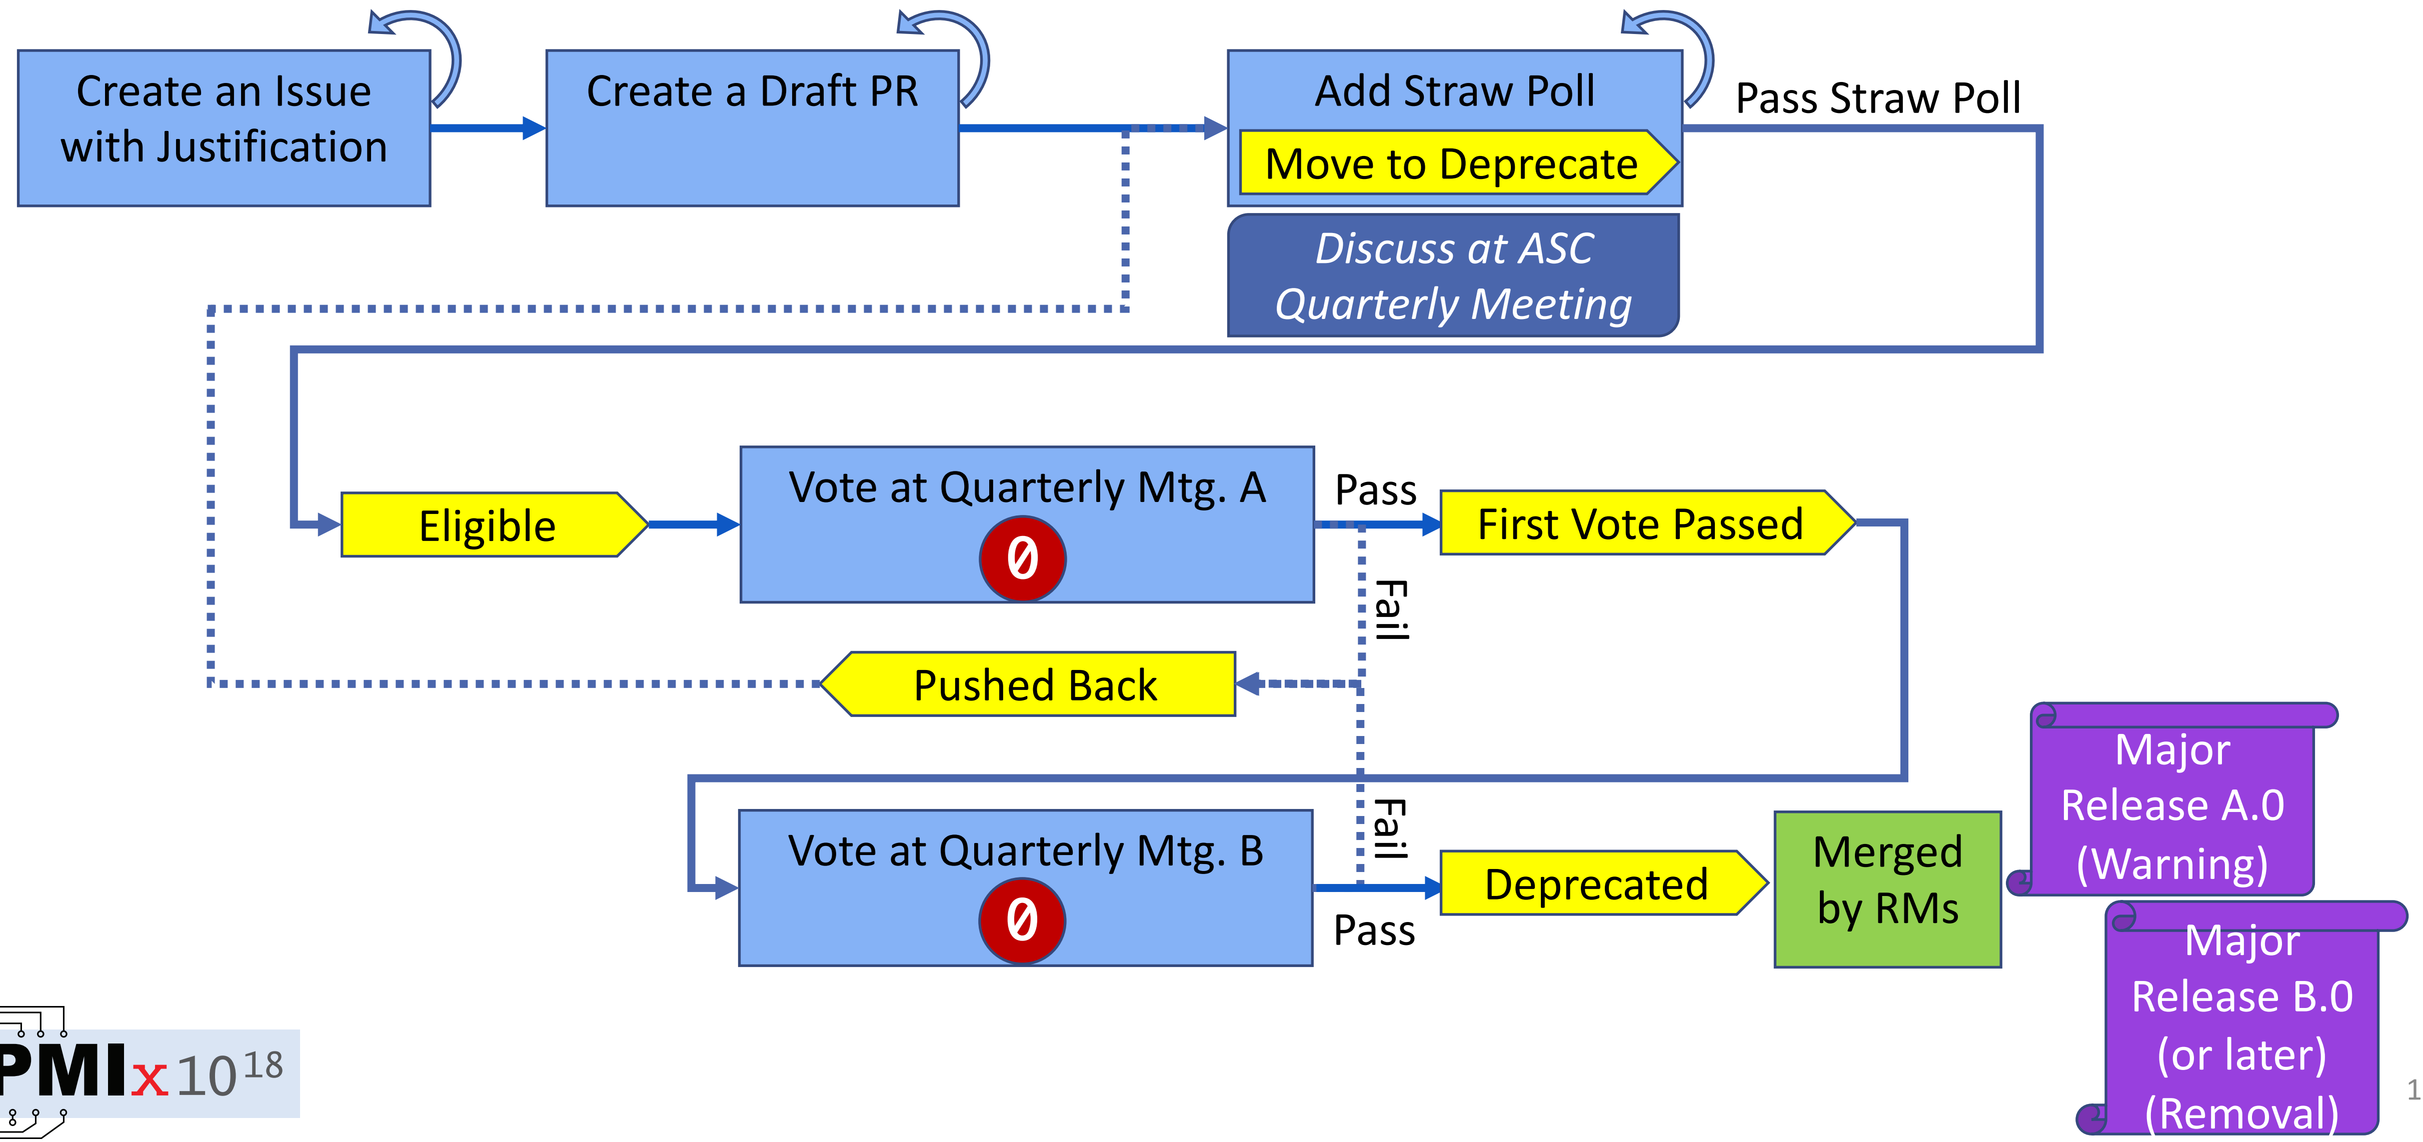

PMIx Standard v5

Transitioning to PMIx Standard v5

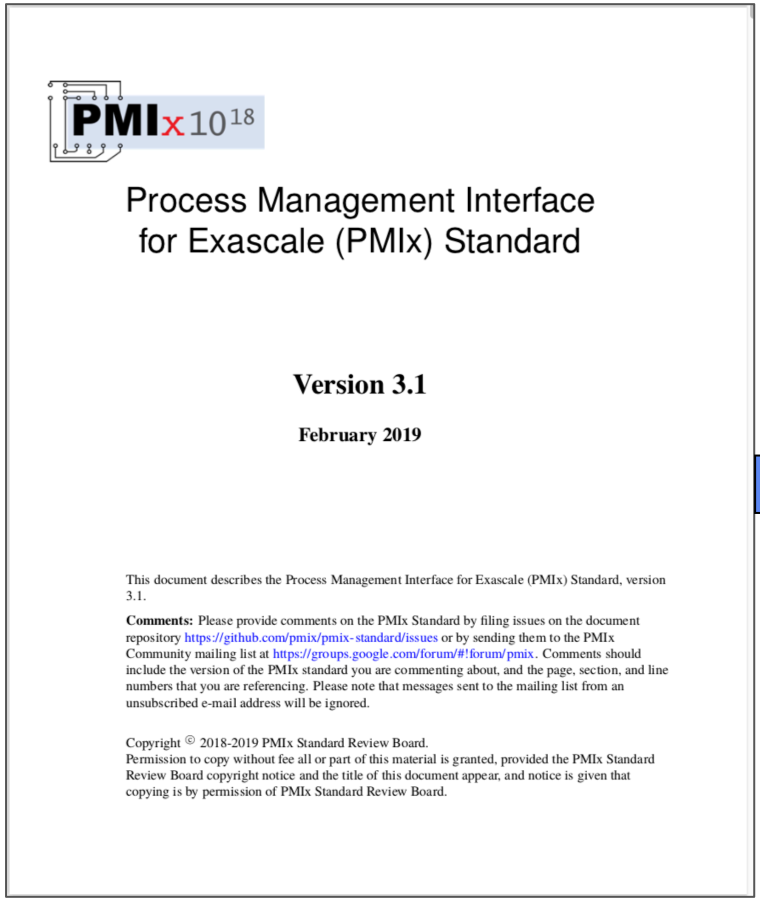

Process Management Interface for Exascale (PMIx) Standard

Version 3.1

February 2019

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 3.1.

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.

Feb. 2019

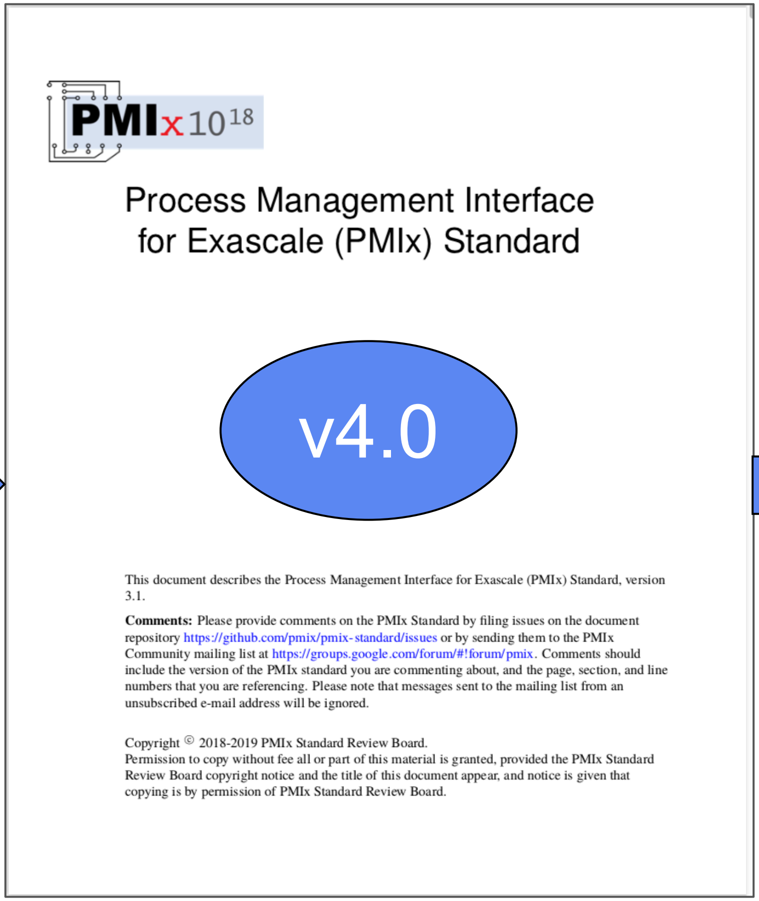

Process Management Interface for Exascale (PMIx) Standard

v4.0

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 3.1.

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 © 2018-2019 PMIx Standard Review Board.
Permission to copy without fee all or part of this material is granted, provided the PMIx Standard Review Board copyright notice and the title of this document appear, and notice is given that copying is by permission of PMIx Standard Review Board.

Dec. 2019

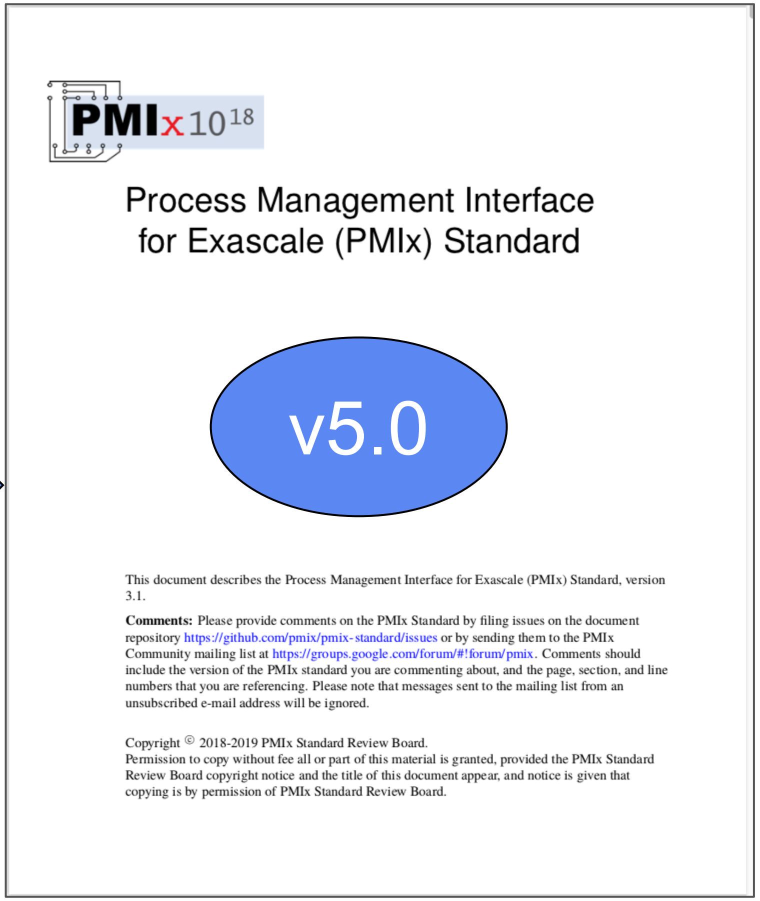

Process Management Interface for Exascale (PMIx) Standard

v5.0

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 3.1.

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 © 2018-2019 PMIx Standard Review Board.
Permission to copy without fee all or part of this material is granted, provided the PMIx Standard Review Board copyright notice and the title of this document appear, and notice is given that copying is by permission of PMIx Standard Review Board.

~2H2020

PMIx Standard v4.0

- **Release Managers:**
  - Ralph Castain, Intel
  - Joshua Hursey, IBM
- **Release Timeline**: Dec. 2019
  - Waiting on descriptive text around tool interactions.
- **Scope of expected changes**:
  - Attribute queries
  - Process groups/sets
  - Network coordinates/topology
  - Extended tool/debugger support
  - Scheduler access to network fabric information
  - Python bindings

PMIx Standard v5.0

- **Release Managers**: TBD
- **Release Timeline**: TBD
  - General guidance is that once any change is voted in then a release timeline will be set.
- **Scope of expected changes**:
  - Re-working of text to become more implementation agnostic
  - Re-organization to better identify slices of functionality
  - Possible introduction of use cases
  - Storage related directives
  - …

Call for ASC Membership

ASC Membership

- IBM
  - Joshua Hursey (Rep.)
  - David Solt (Alt.)
- Mellanox
  - Artem Polyakov (Rep.)
  - Joshua Ladd (Alt.)
- Altair
  - Michael Karo (Rep.)
- OSU
  - Hari Subramoni (Rep.)
  - Karthik Manian (Alt.)
- Intel
  - Ralph Castain (Rep.)
  - Danielle Sikich (Alt.)
- INRIA
  - Brice Goglin (Rep.)
  - Guillaume Mercier (Alt.)
- TU Munich
  - Martin Schulz (Rep.)
- Microsoft
  - Jithin Jose (Rep.)
- Argonne
  - Ken Raffenetti (Rep.)
- Rogue Wave
  - John DelSignore (Rep.)
- LLNL
  - Kathryn Mohror (Rep.)
  - Stephen Herbein (Alt.)
- ORNL
  - Thomas Naughton (Rep.)
  - Swaroop Pophale (Alt.)
- UTK
  - Aurelien Bouteiller (Rep.)

PMIx10^{18}

<Break>

Return at 12:30 pm (US Central)

ASC Leadership Elections

PMIx ASC Leadership Positions – Voting

- **Co-Chair (even-year) 2020-2022** (start Q1 2020)
  - Kathryn Mohror, LLNL
- **Secretary (odd-year) 2019-2021** (start Q1 2019)
  - Stephen Herbein, LLNL
- **Secretary (even-year) 2020-2022** (start Q1 2020)
  - Thomas Naughton, ORNL
- **Release Managers for PMIx Standard v5** (2 positions)
  - Ken Raffenetti, ANL
  - David Solt, IBM

Working Group:
Implementation agnostic document

David Solt (IBM)

- **2020 ASC Meetings**
  - **Q1: January** Propose: Thursday, Jan. 30
  - **Q2: April** Propose: Thursday, April 23
  - **Q3: July** Propose: Thursday, July 16 (face-to-face)
  - **Q4: October** Propose: Thursday, Oct. 15

<Break>

Return at 2:15 pm (US Central)

Dates for 2020 Quarterly Meetings

PMIx ASC Quarterly Meeting Schedule for 2020

The ASC meets quarterly, with at least one quarterly meeting each year to be held “in-person” at an agreed-upon location unless the ASC Members vote otherwise.

ASC quarterly meetings are conducted once per calendar quarter with a preference for a time during the first month of the quarter.

Subsequent changes to a given meeting’s dates must be determined no later than two quarterly meetings before the one being changed.

- **2020**
  - **Q1: January** Propose: Thursday, Jan. 30
  - **Q2: April** Propose: Thursday, April 23
  - **Q3: July** Propose: Thursday, July 16 (face-to-face)
  - **Q4: October** Propose: Thursday, Oct. 15

Governance PRs for Vote

Discuss: ASC Operating Procedures

- Issue: https://github.com/pmix/governance/issues/6
- Pull Request: https://github.com/pmix/governance/pull/7

Conflicting conditions for elevation to Stable status

- Issue: https://github.com/pmix/governance/issues/5
- Pull Request: https://github.com/pmix/governance/pull/4

PMIx 10^{18}

Working Group: Slicing/Grouping of functionality

Stephen Herbein (LLNL)

Discussion Items

Discussion Items

- Open call for new Working Groups
  - Dynamic Workflows
  - Storage
  - ???
- Discussion: PMIx Standard weekly teleconferences
  - Do we need to do these every week or move to every-other week?
  - What should these meetings focus on to help drive progress?
- Discussion: ASC standing agenda item for technical presentations

Thank you for participating today!

SC’19 BoF: Charting the PMIx Roadmap

Date: Wed., Nov 20, 12:15pm-1:15pm

Room: 210-212