

# Designing a Process for Micro-credentials at Scale



## Table of Contents

<b>Table of Contents .....</b>	<b>2</b>
<b>Introduction .....</b>	<b>1</b>
<b>Workflow Overview.....</b>	<b>2</b>
<i>Figure 1: Workflow Overview .....</i>	<i>2</i>
<b>The Details .....</b>	<b>4</b>
<b>Step 1: Create and Validate the Micro-credential .....</b>	<b>4</b>
<i>Figure 2: Issuer Workflow Step 1 .....</i>	<i>5</i>
<b>Step 2: Create Training Benchmarks Database .....</b>	<b>7</b>
<i>Figure 3: Issuer Workflow Step 2 .....</i>	<i>7</i>
<b>Step 3A: Create Training Module.....</b>	<b>8</b>
<b>Step 3B: Build Reviewer Pool.....</b>	<b>9</b>
<i>Figure 4: Issuer Workflow Step 3 .....</i>	<i>10</i>
<b>Step 4: Issue Micro-credentials at Scale .....</b>	<b>11</b>
<i>Figure 5: Issuer Workflow Step 4 .....</i>	<i>12</i>
<b>Step 5: Maintain Reviewer Pool.....</b>	<b>13</b>
<i>Figure 6: Issuer Workflow Step 5 .....</i>	<i>14</i>
<b>Additional Workflow Perspectives: Earner, Reviewer, Recognizer .....</b>	<b>15</b>
<i>Figure 7: Earner Workflow.....</i>	<i>15</i>
<i>Figure 8: Reviewer Workflow.....</i>	<i>15</i>
<i>Figure 9: Recognizer Workflow.....</i>	<i>16</i>
<b>Topics for Further Exploration .....</b>	<b>18</b>

This paper was produced by Michelle Riconscente, Designs for Learning, Inc., in collaboration with Karen Cator, Jennifer Kabaker, Odelia Younge, and Robert Bajor. The author would like to thank Barnett Berry, Dan Brown, Ann Byrd, Alesha Daughtrey, Pat Deklotz, Lisa Dieker, and Brent Maddin, as well as participants in the August 2016 Digital Promise Micro-credential meeting, for their helpful and informative input on the workflows.

This publication was made possible in part by a grant from Carnegie Corporation of New York. The statements made and views expressed are solely the responsibility of the authors.

## Introduction

Micro-credentials represent a promising new direction for transforming educator professional development into effective, personalized learning opportunities with direct impacts on classroom instruction. There are many components to a fully functioning ecosystem of micro-credentials, including a robust library of learning opportunities, systems for creating, accessing, and earning them, as well as viable business models and incentives for participation. In 2013, Digital Promise embarked on a journey to develop micro-credentials with relevance and appeal to educators; design technology systems to support the logistics of creating, earning, issuing and sharing micro-credentials; and pilot their use in a variety of contexts, from university pre-service programs to district and state initiatives.

Numerous partnerships now power this effort, including collaborations with over thirty issuing organizations that have collectively developed over 250 micro-credentials in a variety of content areas. Additionally, states and districts are recognizing earned micro-credentials by awarding continuing education units (CEUs), pay scale advances or bonuses. At least one state is piloting the use of micro-credentials towards re-licensure.

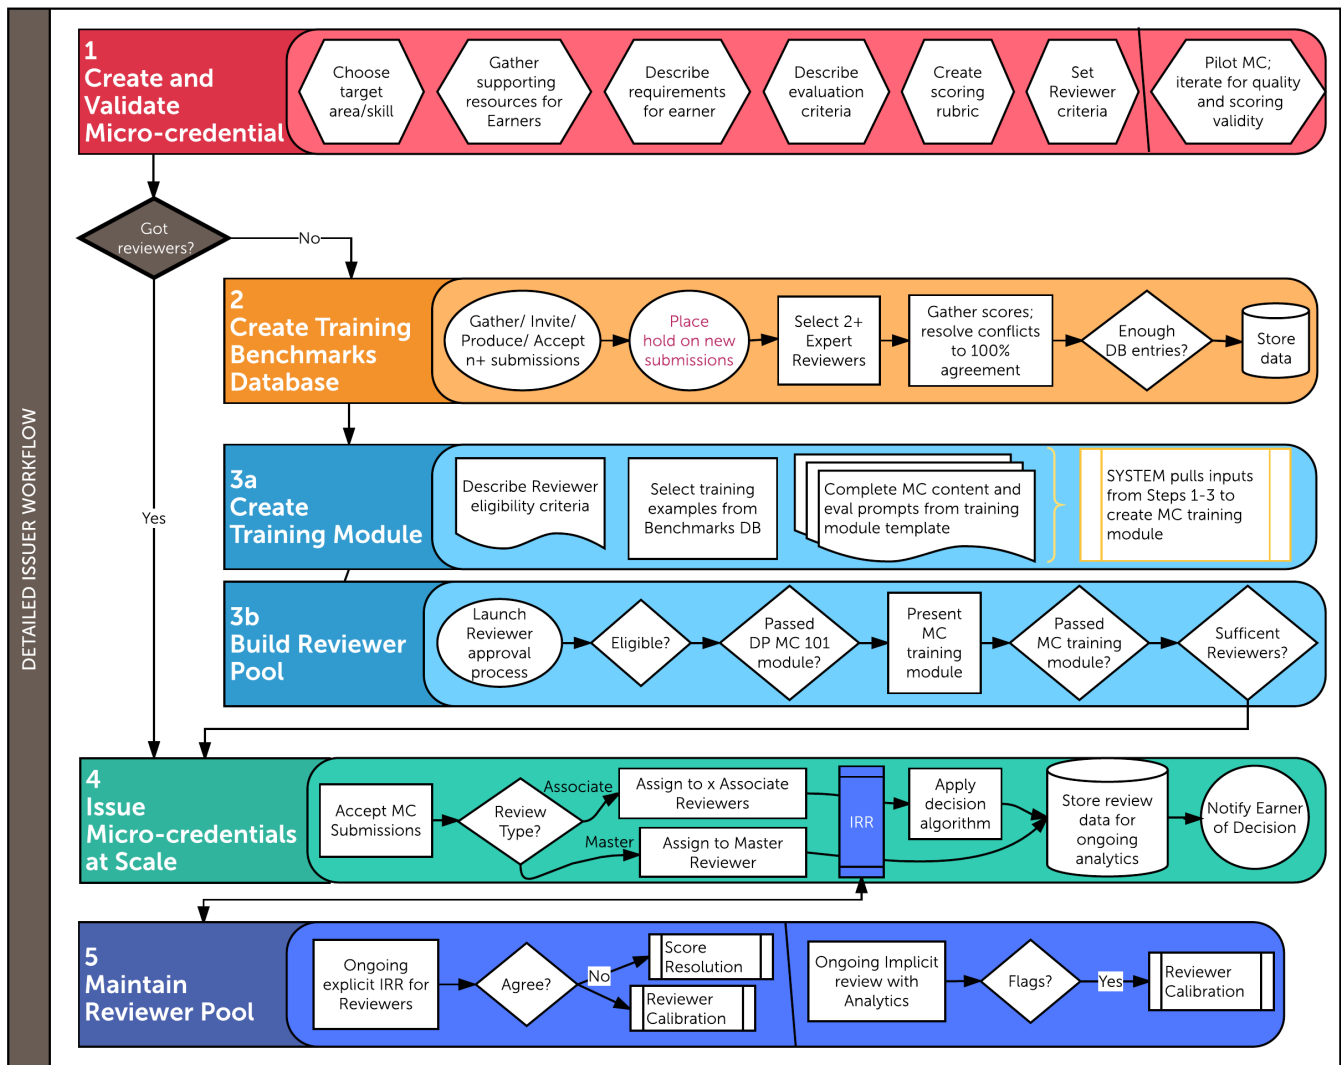
As the potential of micro-credentials is increasingly recognized and interest grows in broad adoption, it has become important to design strategies for large-scale implementation. In many cases, the volume of submissions will outgrow the issuing organization's bandwidth for evaluating them. Consequently, critical to the success of micro-credentials at scale is a review process capable of handling large numbers of submissions in a timely fashion, while maintaining a high level of reliability and validity of the evaluation process.

This guide presents a solution to the challenge of scaling the review process. It lays out a workflow for creating a pool of qualified reviewers, and for maintaining consistency of the review process over time. A driving consideration in the development of these flows was the importance of building the credibility of micro-credentials. Describing a common process and providing transparency into the processes by which micro-credentials are scored and reviewers are certified is key to gaining trust of key stakeholders. Stakeholders include educators who would earn them, as well as the districts, states and organizations who will recognize them. The following section provides a high-level overview of the five steps and subsequent sections walk through the workflow in detail.

## Workflow Overview

An overview of the entire workflow is presented in Figure 1, and covers the micro-credential trajectory from creation through long-term maintenance at scale. The first step is the creation of the micro-credential itself, including targeted competencies, submission requirements, evaluation criteria, and scoring rubrics. The validation of the micro-credential also takes place in this first step, to confirm that those who earn the micro-credential are indeed demonstrating the competency as described, and also to ensure that the instructions and processes associated with the micro-credential are operating as intended.

Figure 1: Workflow Overview



Following Step 1, a key branch point – “Got Reviewers?” – appears in the flow. For most organizations, a successful implementation would generate more submissions than can be reasonably reviewed by staff. Steps 2–3 describe a template-driven process for building a training program and certifying individuals who successfully complete the training to serve as reviewers. The central approach to reviewer training and certification is based on their ability to score a representative set of submissions on par with issuer-determined submission scores. In Step 2, the issuer creates a Training Benchmarks Database. For each micro-credential, the issuer would be responsible for creating a library of expert-scored submissions—a “Benchmarks Database”—comprising submission examples reflecting each score option. The minimum size of the Benchmark Database is determined by a formula based on the scoring options, expectations for the number of attempts a potential reviewer can make, and item security considerations.

In Step 3, the issuer creates a training module and builds a reviewer pool. To ensure quality and consistent reviews of submissions, each micro-credential to be scaled would have a training module associated with it. Training modules would follow a common template, both to expedite their creation, and to provide consistency across the library of micro-credentials. Once the Benchmarks Database and training materials are in place, the system opens to reviewer applicants. When a large enough pool of qualified reviewers is in place, the system activates the micro-credential for new submissions.

Step 4 is where micro-credentials are issued at scale. With a system overseeing the submission and review process, micro-credentials can be much more easily and efficiently issued at scale. As micro-credential seekers submit their materials, the system batches them, assigns them to appropriate reviewers, and ushers submissions through the review process. Discrepancies in reviewer scores trigger flags to the issuer and a process for score resolution. Earners are automatically notified of the outcome of their submission and encouraged either to resubmit, or to pursue related micro-credentials or to complete a related stack.

Ongoing maintenance of the reviewer pool takes place in Step 5. Maintaining a quality reviewer pool is essential to the ongoing validity and credibility of micro-credentials. Step 5 is an ongoing process with two strands. The first strand is the explicit maintenance. Reviewers are periodically assigned benchmark samples to score, or paired with additional reviewers to check for inter-rater reliability. Meanwhile, ongoing implicit maintenance is carried out using system analytics, such as whether a reviewer tends to score higher or lower than others.

## The Details

### Step 1: Create and Validate the Micro-credential

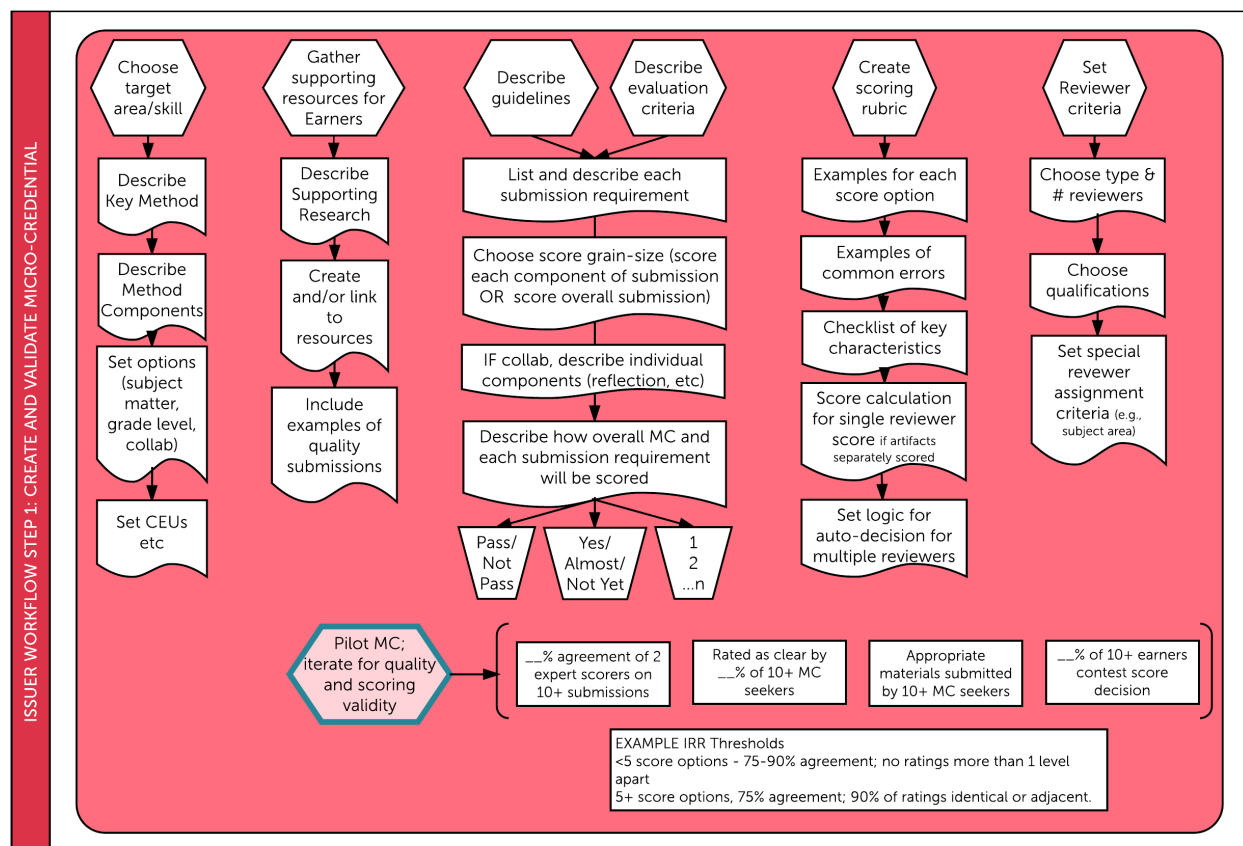
The first step is the creation of the micro-credential itself. Most components of this process are already familiar to issuers, and the design shown in Figure 2 follows the existing Digital Promise workflow for submitting micro-credentials to the platform. The issuer describes the Key Method – for instance, crafting driving questions – as well as the Method Components, outlining what it means to implement this competency in practice, and why it matters. In addition, looking ahead to an automated review system, the issuer could also decide whether or not to allow collaborative submissions, and set options to enable earners to indicate their domain (ELA, science, etc.). When defining the micro-credential, the system may also prompt the issuer to measures of difficulty such as number of equivalent Continuing Education Units (CEU's).

Although including options such as collaboration or target domain is not necessary in a minimally viable system for supporting micro-credentials, capturing these features as part of the creation process serves two purposes. First, it surfaces questions at the outset of the design process for issuers to consider, encouraging thoughtful decision-making on scenarios they are likely to encounter once the micro-credential is launched. Second, capturing these design decisions and characteristics in the platform makes it possible to automate related processes, such as batching collaborative submissions from members of the same team, distinguishing collaboration from plagiarism, and assigning reviewers with domain-specific knowledge to submissions. From a data perspective, explicitly indicating these characteristics of micro-credentials enables a variety of research and analytics activities.

As shown in Figure 2, in creating the micro-credential, an issuer will gather supporting materials for earners, including relevant research, links to resources, and examples of quality submissions. The micro-credential also needs to describe submission requirements, the scoring scale, and the grain-size for scoring. The scoring scale refers to the possible outcomes of the micro-credential evaluation process. For example, the scale might include pass/resubmit; pass/almost there/not yet; high pass/pass/almost there/not yet. In this first step, issuers also decide whether each submission element will be scored separately, with a final score calculated from the individual scores, or whether the micro-credential will be scored holistically. For collaborative submissions, the issuer may opt to require an additional individual component such as a personal reflection, or an explanation of the specific contributions s/he made to the group submission. Part of creating a new micro-credential also includes generating a scoring rubric that reviewers will apply to determine a score. Ideally a rubric includes examples of each scoring option, as well as illustrating common

errors to make it easier for reviewers to quickly detect them, and potentially a checklist of key characteristics to look for.

Figure 2: Issuer Workflow Step 1



To enable automated processes down the line, the system would also prompt issuers to indicate how outcome decisions will be made. If each part of the micro-credential submission will be scored individually – for example, if the main description, the explanation of student work, and the reflection each get a separate score – the system could offer the issuer a variety of options for how to combine those scores into a final micro-credential outcome. These options could be a simple average, threshold requirements for each section, or weighted sum. At this point, the issuer could also describe the qualifications of reviewers who could be approved to score submissions. These might include having earned the micro-credential themselves, having specific subject area expertise, or other credentials such as National Board Certification.

The issuer would also decide how many reviewers should score each submission, and what logic should be applied when the reviewer scores do not agree. The workflows illustrate two general reviewer approaches: a master and associate. Master reviewers function in a context in which a single reviewer makes a micro-credential score determination. This model is akin to pre-service programs in which a faculty member assigns a grade to a learner's coursework, and is typically grounded in a reviewer pool of highly qualified and vetted experts. However, in many cases, scaling a micro-credential assessment workflow will demand the engagement of individuals previously unknown to the issuer organization.

For this latter scenario, the workflows describe an associate reviewer process in which multiple reviewers are assigned to each submission, with the principle of convergence providing confidence in the outcome decision. Including multiple reviewers is analogous to adding legs to a table – the more legs, the more stable the table becomes. Moreover, the data captured by multiple reviewers offer insight into the clarity of the micro-credential itself, as well as an opportunity to ensure that individual reviewers tend to apply similar judgment in applying the scoring rubric to submissions.

As show in the lower portion of Figure 2, the step for creating the micro-credential includes a piloting and validation process for confirming that the instructions and scoring are trustworthy and reliable. The pilot and validation process will range considerably in scope and depth, reflecting the fact that micro-credentials span a broad range of competencies, and vary in their scoring complexity. Whereas one might focus on data analysis and require responses to a set of questions whose answers are clearly right or wrong, another may target practices that are more challenging to assess, such as a change in mindset. Therefore, the complexity of the micro-credential, and in particular of the scoring process, will play a role in determining the complexity of the validation process.

At a minimum, a validation process could require high rate of agreement between two expert scorers on ten or more submissions. Issuers could set the threshold for agreement depending on the stakes and the scoring scale. For example, if a micro-credential were scored on a 4-point scale, the issuer might not require perfect agreement between "high pass" and "pass" but might require perfect agreement on "pass/not yet" outcomes. In addition to checking for scoring consistency, other considerations in validating the micro-credential include ensuring that the instructions can be easily understood by earners, by checking that earners are responding to the submission requirements as intended by issuers, and that outcome decisions are not frequently contested by those who sought to earn the micro-credential. Addressing these latter considerations becomes all the more important in a large-scale system, so that implementation runs as smoothly as possible.



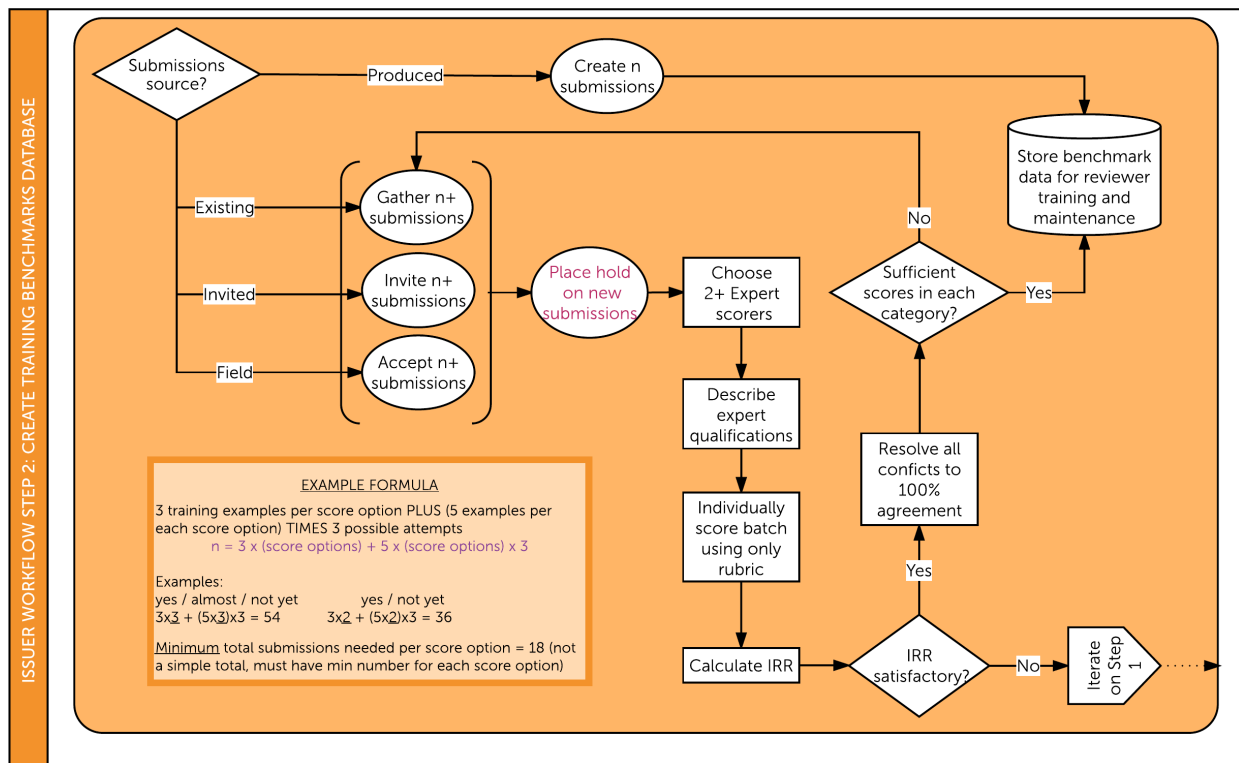
A variety of approaches can be taken to validate micro-credentials, and these processes would be carried out at the discretion of the issuer. By clearly describing the validation approach taken, those considering earning or recognizing the micro-credential will be better informed. Transparency around these processes, drawn when possible from a common set of options, will build trust in the validity and credibility of the micro-credentials.

## Step 2: Create Training Benchmarks Database

A key feature of the reviewer scaling process described in this document is the use of scored benchmarks within the reviewer training and maintenance process. For each micro-credential, the issuer would be responsible for creating a set of expert-scored submissions, representing each score option.

As shown in Figure 3, there are a variety of sources an issuer could turn to for creating this Benchmarks Database. For issuers who have previously begun receiving submissions, the existing library of submitted work could serve this purpose, provided there are sufficient examples in each score category. Issuers could also produce their own examples of submissions corresponding to each level of the scoring rubric, however it is recommended that authentic examples be used when possible given their increased likelihood of reflecting the kinds of submissions that will be received.

Figure 3: Issuer Workflow Step 2



The call-out box in Figure 3 presents a formula for determining the number of examples required for the database. The minimum size of the benchmark database is determined by a formula based on the scoring options and expectations for the number of attempts a potential reviewer can make in becoming certified as a reviewer. The formula assumes that the training module for reviewers will include three examples for each score option. This assumption is reflected in the first part of the equation: “3 x (score options)”. The second assumption is that prospective reviewers would have up to three attempts at passing the training module, and would have to accurately score up to five examples at each score option in order to be certified.

Expert scorers, identified by the issuer, are key to this process, since the scores of submissions in the database will serve as “ground truth” for the reviewer scaling process. As the workflow shows, at least two experts would independently score each submission using the rubric—in other words, applying the same process that reviewers will use to score submissions. Inter-rater reliability would be calculated as a check on the scoring process. If it turns out that experts have a hard time agreeing on submission scores, the process would return to Step 1 to revisit the scoring rubrics and requirements. However, assuming that acceptable inter-rater reliability is achieved, any conflicts would be resolved in conference by the experts such that one final score is assigned to each database entry. At this point, the system would check to see whether sufficient entries had been found for each score option. If so, this part of the process is complete. Otherwise, the system would loop back to generating or identifying additional submissions to include in the database.

It is worth noting that as the adoption of micro-credentials increases over time, considerations related to security of the training database will need to be addressed. In large-scale test development systems, ensuring security often means having such a large number of examples that it becomes difficult to game the system by knowing what all the test cases will be. Another approach, which would apply to some but not all micro-credentials, is dynamic setting of variables within submissions, similar to varying math problems which have the same structure but differ in the particular numbers presented. The logic informing the formula presented here is designed to yield a micro-credentials Benchmarks Database for which it is unlikely to have dishonest reviewers.

### **Step 3A: Create Training Module**

To ensure quality and consistent reviews of submissions, each micro-credential to be scaled would have a training module associated with it. Training modules would follow a template both to expedite their creation, and to provide consistency across the library of micro-credentials. A typical training module creation flow includes a description of reviewer eligibility criteria, and a selection of three examples at each score level from the Benchmarks Database created in Step 2.

As shown in Figure 4, issuers would complete standard fields regarding the rationale for the micro-credential, and an overview of the requirements and materials that earners can access as they prepare the process of creating their submission for the micro-credential. Central to the training module is the presentation and description of the scoring rubric, together with its annotated application to the three examples from each score category, pointing out key characteristics as well as common errors. For micro-credentials that offer feedback in addition to the final outcome of the scoring process, guidance would be provided together with examples of both effective and ineffective feedback.

Implementation of the training module could take several forms. Through a common platform to minimize issuer time investment, the system could pull responses to standard prompts into a consistent training module format. For example, the issuer could simply select items from the Benchmarks Database to use in the training module, and those examples could be pulled in by the system in the process of generating the training module.

An important area to explore is whether a single training module could serve several micro-credentials. If the review process is too work-intensive, it will not be able to draw individuals to become reviewers; at the same time, a high bar must be maintained to ensure reviewer quality and consistent scoring of micro-credentials. As issuers move forward with this component of the reviewer scaling process, the right balance point will become clearer, and will likely need to be adjusted over time as experience and insight grow.

### Step 3B: Build Reviewer Pool

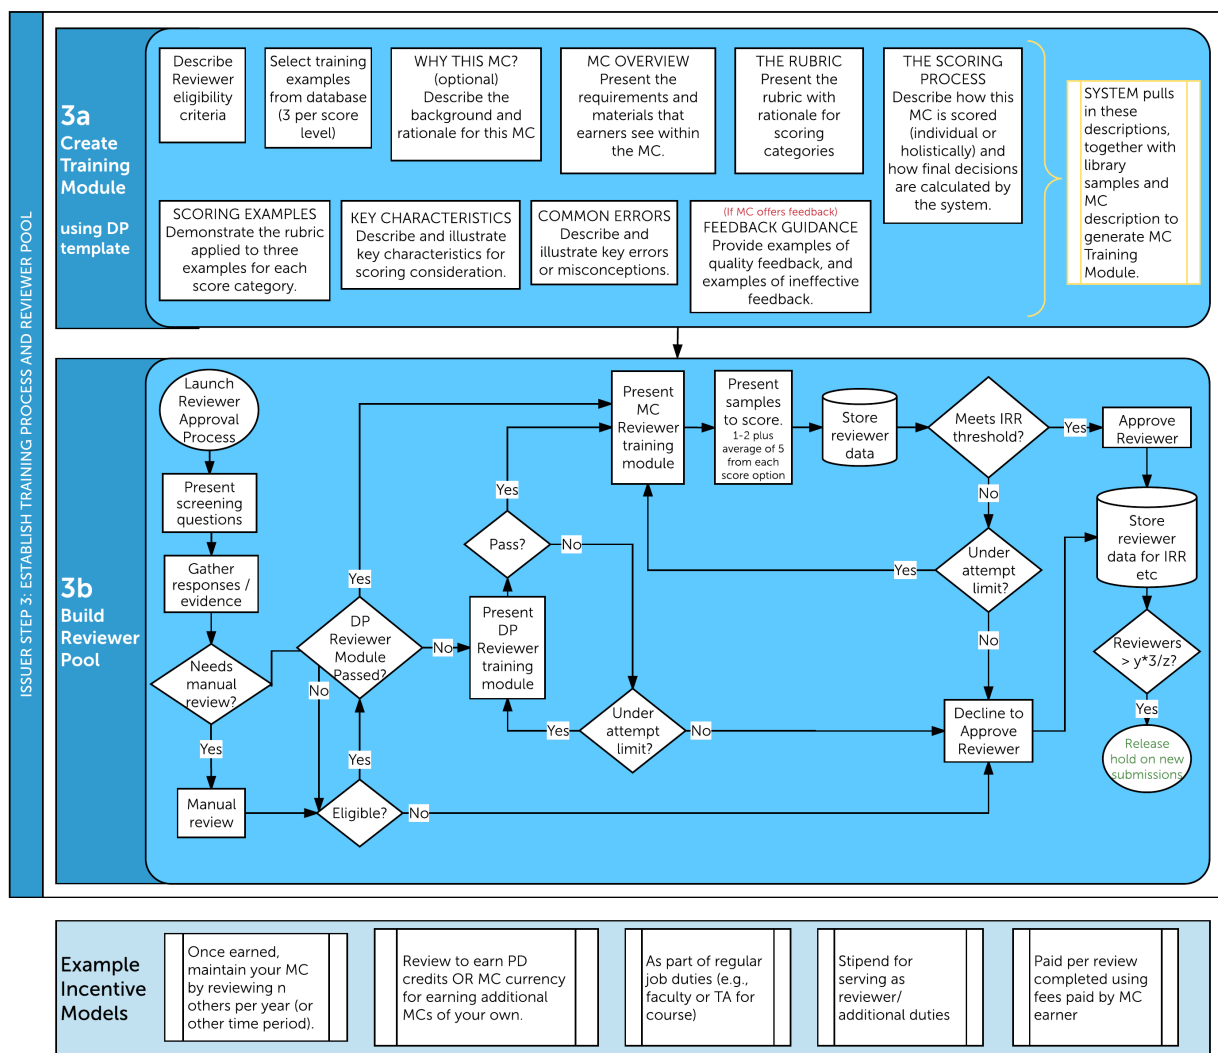
Once the benchmark database and training materials are in place, the system accepts applications for becoming a reviewer. When a sufficiently large pool of qualified reviewers is in place, the system activates the micro-credential for new submissions. The minimum size of the reviewer pool can be determined on the basis of the number of anticipated submissions within a scoring period, the number of associate reviewers who will review each submission, and the number of submissions that will be assigned to each reviewer, as follows:

$$\text{minimum \# reviewers needed} = \frac{\text{anticipated \# of submissions} \times \text{\# of reviews per submission}}{\text{\# of assignments per reviewer}}$$

For example, if the issuer anticipates receiving 200 submissions per month, having three reviewers score each submission, and assigning 10 submissions per month to each reviewer, then the formula would be:  $(200 \times 3) / 10 = 60$  reviewers. The second portion of Figure 4 illustrates the workflow for building the reviewer pool. In a comprehensive platform, this process would be automated by the system, with little involvement required from the issuer.

As a first step, the reviewer applicant would be presented with screening questions based on the eligibility criteria. Some of these may be automatically pulled in by the system, for instance in the case of requiring reviewers to have already earned the focal micro-credential. Others may require manual review, for instance confirmation of employment as a faculty member in a school of education, or certification status such as National Board Certified.

Figure 4: Issuer Workflow Step 3



Reviewer applicants who meet the eligibility criteria would move on to complete a general module created by Digital Promise. This module is completed just once across all micro-credentials, and covers the core elements of the system, including general guidelines for applying scoring rubrics and for providing quality feedback. Having successfully completed

the general training module, the reviewer applicant would then complete the specific micro-credential reviewer training module, including submitting his or her scores of the five samples randomly drawn from the Benchmarks Database for each score option. Some issuers may choose to present more or less than five examples at each score level. The rationale for that number is to balance the workload while making it nearly impossible for an unqualified individual to guess their way to successful completion of the reviewer training module. These probabilities can be objectively calculated in advance, to inform an issuer's decision regarding the number of scored examples and level of agreement. For instance, for a micro-credential that is scored pass/not yet, the chances of guessing with 100% accuracy would be approximately  $0.5^{10} = 0.001$ , or less than 0.1%. The inclusion of quantifiable criteria in the reviewer training process contributes to transparency and rigor, which together serve to build trust and credibility in the broader community.

As for the process of creating the micro-credential itself, issuers maintain control over several levers in setting up the reviewer training module, such as the number of score levels, the number of retry attempts a reviewer applicant has, the amount of time a reviewer applicant must wait between retries, and the level of agreement (i.e., inter-rater reliability) that must be obtained to be approved as a reviewer. At the same time, the common elements of the process described here lend consistency across micro-credentials and support transparency.

Though a full exploration and discussion of incentives for becoming a reviewer is beyond the scope of this paper, there are several promising avenues to pursue. These include models that require individuals to maintain their reviewer status by reviewing a certain number of submissions periodically; the opportunity to earn credits in the system as a reviewer that could be applied to earning additional micro-credentials; stipends for serving as a reviewer; and paid review completed using fees from micro-credential seekers.

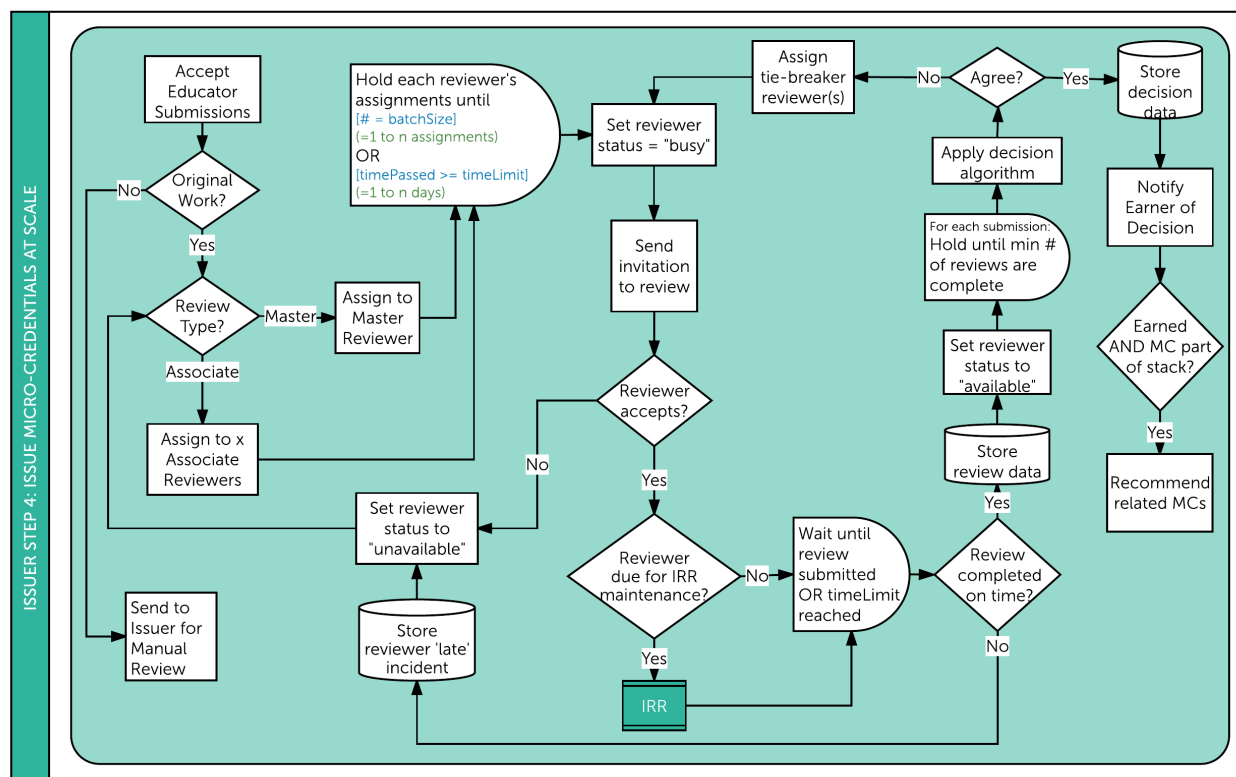
#### **Step 4: Issue Micro-credentials at Scale**

With a system overseeing the submission and review process, micro-credentials can be easily and efficiently issued at scale. As potential micro-credential earners submit their materials, the system batches them, assigns them to appropriate reviewers, and ushers submissions through the review process. Discrepancies in reviewer scores trigger flags to the issuer and a process for score resolution. Earners are automatically notified of the outcome of their submission and encouraged either to resubmit or to pursue related micro-credentials or to complete a related stack.

The workflow presented in Figure 5 is set in motion when a prospective micro-credential earner submits an application through the platform. In a robust system, a potential first step could apply automated text analysis to the submission as a quick check against plagiarized

content. Submissions flagged as suspicious would trigger a notice to the issuer for manual review. Based on the characteristics of the submission and the reviewer requirements set by the issuer, the submission would be assigned to reviewers whose status is 'available' in the system. The number of reviewers per submission is set by the issuer in Step 1. To improve scoring consistency, as well as to streamline the process for reviewers, the flow uses a batching strategy, with the batch size set by the issuer. Once a reviewer has been assigned a certain number of submissions to review—or if a specified period of time has passed—the system sets that reviewer's status to busy and generates an email notice inviting the reviewer to accept the review task. Similar to processes for academic journal reviews, the reviewer has the option to accept or decline the assignment prior to viewing the submissions. This feature would enable the platform to dynamically adjust for periods of reviewer unavailability, due either to vacation or other commitments.

Figure 5: Issuer Workflow Step 4



Upon reviewer acceptance of the assignment, the system checks to see whether explicit IRR maintenance processes should be activated. The flow then pauses until either the reviewer completes the reviews, or until a deadline is reached. The system sends reminder emails as

the deadline approaches to increase reviewer response rates and timely completion of the task. In the event that the time limit is exceeded, the system loops back to assign a replacement reviewer, following the same flow again from invitation through to pausing for reviewer scoring or a time-out. Provided the reviews are completed on time, reviewer status is set back to “available”, possibly with a delay so they have some time off between receiving batches of review assignments.

Once all reviews are complete for a given submission, the system automatically calculates a decision, applying the criteria set by the issuer in Step 1. When reviewers’ scores do not agree, the system automatically pulls in one or more additional reviewers. As the system captures data on reviewer score agreement over time, helpful insights can be generated for issuers, as frequent occurrences of discrepant scores may indicate a need for more clarity in the micro-credential requirements, scoring rubrics, or reviewer training module.

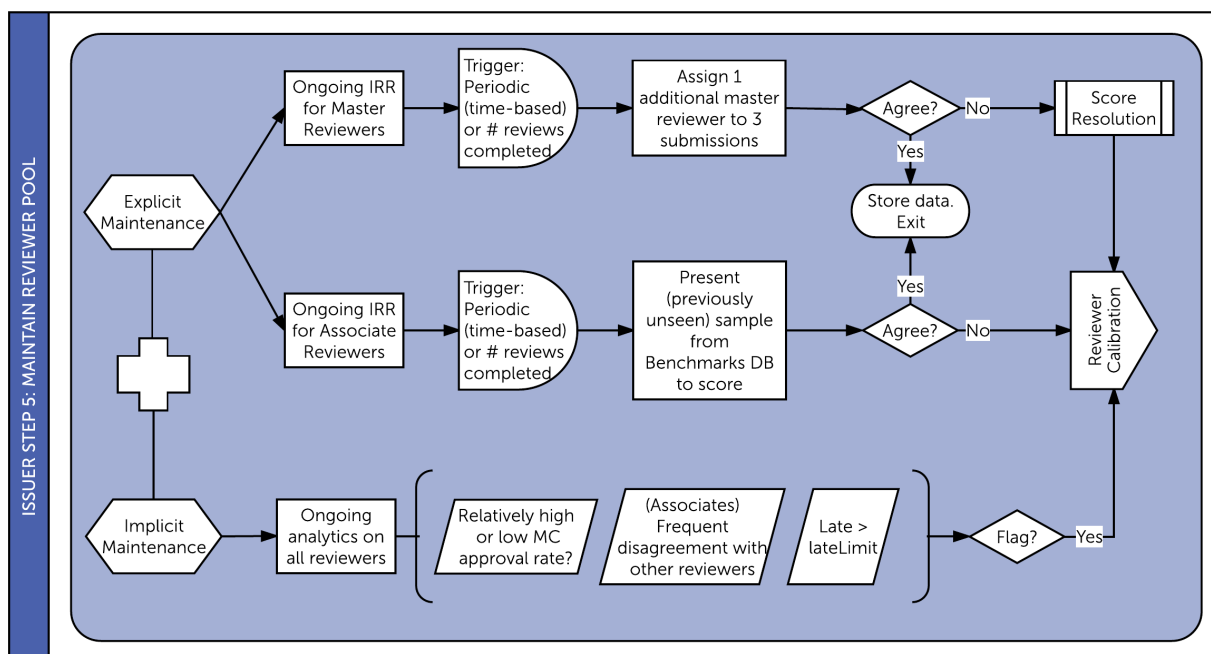
At this point, the earner is notified of the outcome of their submission. In cases where the micro-credential has not been granted, the system might encourage the seeker to resubmit, and there may be a wait time imposed by the system to encourage thoughtful resubmissions.

When the outcome is positive, the system might recommend related micro-credentials, based on related competencies, micro-credentials sought by others who also have the newly earned micro-credential, or micro-credentials that belong to the same stack as the newly earned one. If the earner now qualifies to become a reviewer of the micro-credential, the system could offer congratulations and an invitation to become a reviewer.

### **Step 5: Maintain Reviewer Pool**

Maintaining a quality reviewer pool is essential to the ongoing validity and credibility of micro-credentials. Step 5 is an ongoing process with two strands, labeled “explicit” and “implicit” in the workflow shown in Figure 6. Over time, it is natural for trends in score outcomes to shift, such that reviewers become either more or less stringent in applying rubrics. A robust system monitors these trends to maintain consistency of scoring.

Figure 6: Issuer Workflow Step 5



In explicit maintenance, reviewers are periodically assigned benchmark samples to score, or paired with additional reviewers to check for inter-rater reliability. In the case of micro-credentials that are reviewed by single master reviewers, the system could periodically assign the same submission to two master reviewers, and confirm that the same score was awarded by each. This process is invisible to the reviewers during the review process. When discrepancies are detected, the system would notify the issuer, and depending on issuer settings, might redirect a reviewer to complete the training module again. The reviewer could be temporarily set to “unavailable” so that they do not receive additional assignments until they have re-demonstrated acceptable inter-rater reliability. Alternatively, master and associate reviewers may be periodically assigned an entry from the Benchmarks Database to score, to confirm that they are still applying the scoring rubric consistent to issuer expectations.

Meanwhile, ongoing implicit maintenance is carried out using system analytics, such as whether a reviewer tends to score higher or lower than others across multiple submissions. Discrepancies occurring frequently for a specific micro-credential would indicate potential issues with the submission requirements, resources, scoring rubric, or training module. Data would also flag specific reviewers who tend to apply stricter or more generous scores on specific micro-credentials relative to the reviewer pool. These patterns would trigger notices to the Issuer, and potentially also result in temporarily removing a reviewer from active duty until improved scoring consistency with the Benchmarks Database is demonstrated.



## Additional Workflow Perspectives: Earner, Reviewer, Recognizer

The workflows described in this paper focus on the perspective of issuers of micro-credentials. Micro-credential earners, those who review submissions and those recognize them are additional important roles in a complete ecosystem. Figures 7-9 sketch some preliminary workflows from those perspectives.

Figure 7: Earner Workflow

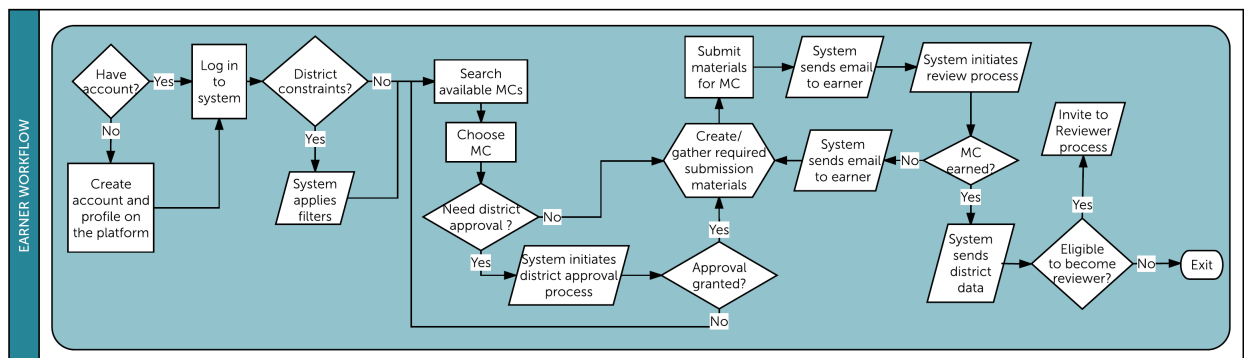


Figure 8: Reviewer Workflow

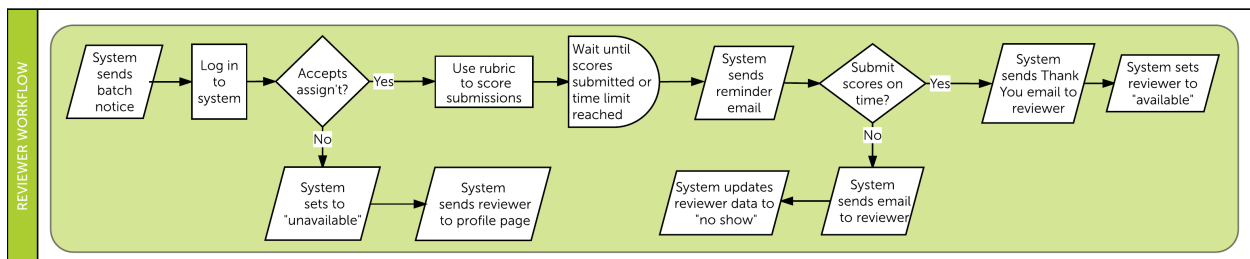
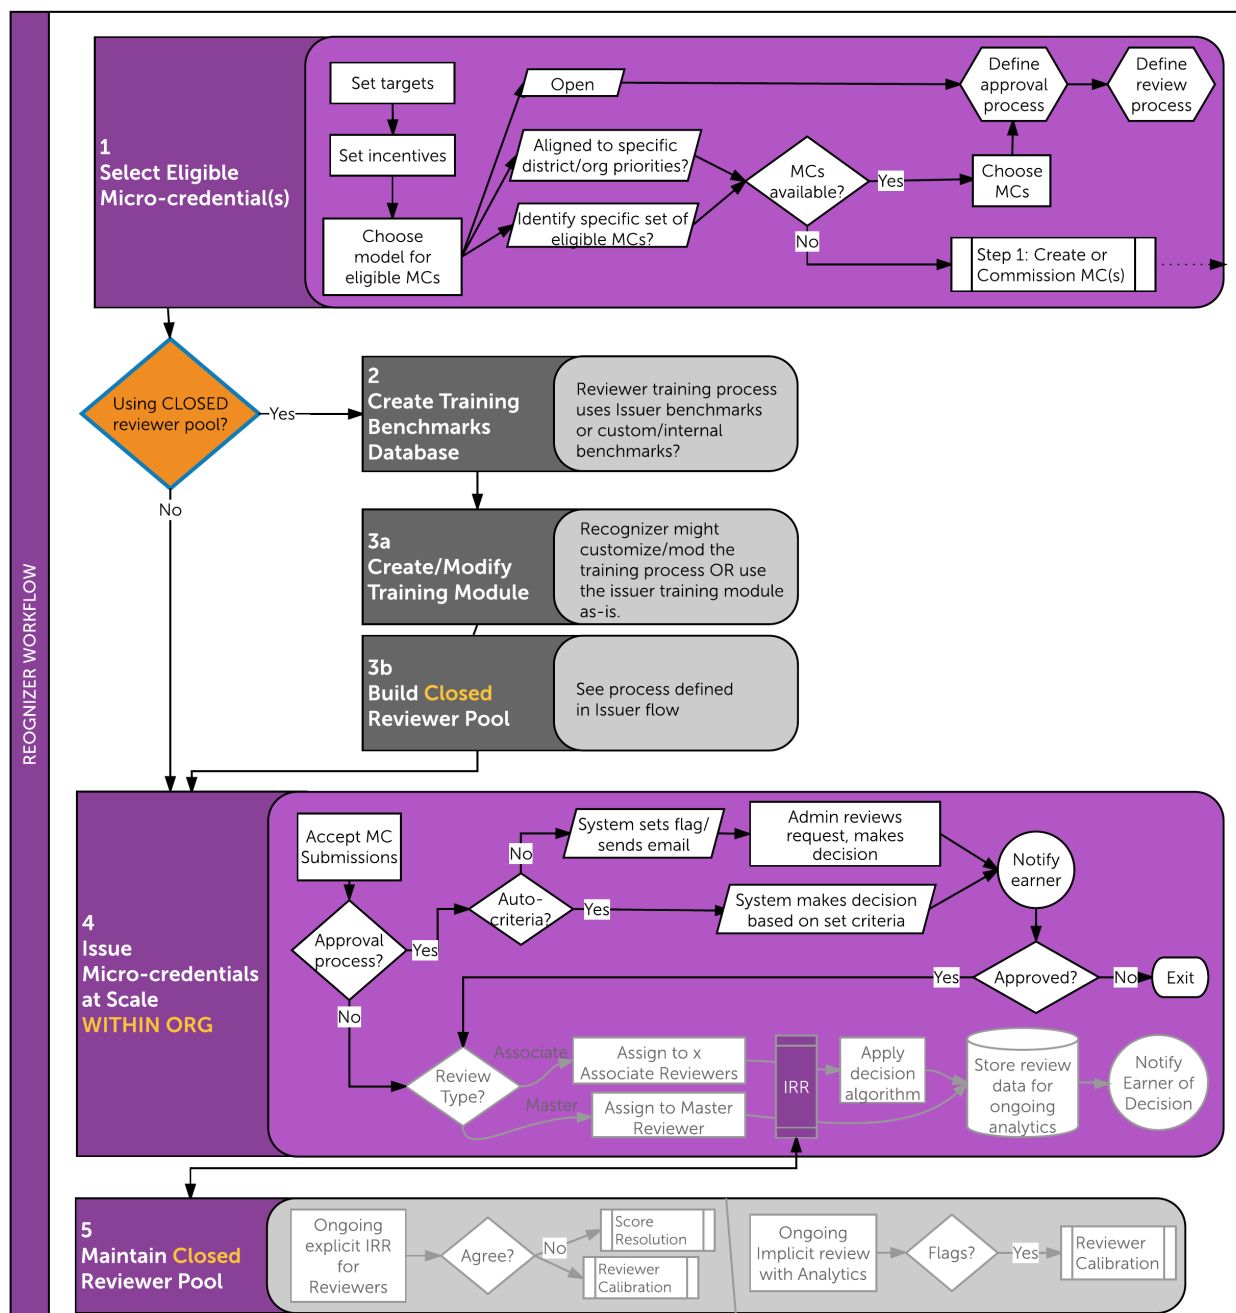


Figure 9: Recognizer Workflow



For micro-credential earners, the core process remains fairly straightforward. For those earning micro-credentials within the context of a specific district or state program of professional development, relevant micro-credential filters or approval processes may be included in addition to selecting and preparing their submission materials. For reviewers in the system, the process parallels the Issuer workflows.

Future design work could explore approaches to surfacing the maintenance of score consistency over time. For example: Would reviewers have access to a dashboard showing the level of agreement with other reviewers? Would they be notified at the same time as Issuers when patterns of inconsistency emerge in the data? Would they be notified when they had successfully passed a check on inter-rater reliability? A robust system would also include vehicles for receiving reviewer input and feedback on the process, including the scoring rubrics and the review process itself.

The workflow for recognizers of micro-credentials becomes potentially more complex, depending on the context of adoption. The main use case shown in Figure 9 reflects a district perspective. Step 1 walks through the process of determining which competencies and purposes micro-credentials would serve in the broader district ecosystem, and whether micro-credentials already exist to address those purposes or need to be created. Additionally, relevant micro-credentials might be selected from the entire library or a district filter could be applied. The district may also craft a process for requesting approval to pursue micro-credentials.

In some cases, a district may choose to create its own reviewer pool for a set of micro-credentials, either out of an interest in overseeing the review process, to guarantee reviewers are available for the volume of submissions the district will generate or to better control the costs of the process. In this scenario, the process would closely follow the one described in the Issuer workflow, creating a closed reviewer pool that uses the same or modified training modules.

This section discussed a sampling of considerations related to key roles in a broader ecosystem. As the adoption of micro-credentials evolves, a variety of use cases and innovations will follow, requiring modifications and enhancements to the processes detailed here, and opening new views to the potential of micro-credentials to serve and advance professional learning in the education sector.

## Topics for Further Exploration

The workflows described in this paper represent a sampling of many possible strategies for scaling the review process in a robust micro-credential ecosystem. Even with the detail provided here, several topics remain to be explored and addressed for various use cases. This section offers a brief discussion of some of these topics for further consideration and iteration.

Implementing a scaled review process presupposes that a broader context in which micro-credentials are valued by key stakeholders across the education ecosystem—from educators, to school districts and learning organizations. The work of Digital Promise to convene a community of issuers and recognizers, curate a library of quality micro-credentials, and build practical tools to enable their implementation, is already moving the field ahead in this regard. Putting in place transparent and consistent review processes on par with the criteria for micro-credentials themselves can contribute to the trust-building and credibility essential for long-term effectiveness of a micro-credentialing approach to effective professional learning.

The role of feedback in the micro-credential review process is an area in need of additional exploration. Issuers' reported experiences with scoring micro-credentials suggests that with modest training, most reviewers successfully attain high inter-rater reliability rates. However, reviewers' feedback on submissions tends to be more varied both in focus and quality, and may require additional investment of time to develop the relevant benchmarks and for adequate training. While feedback can offer earners richer opportunities for growth and learning, variability in feedback received could be perceived by earners and other stakeholders as undermining the scoring decisions. One possible solution is to create a training module with rigorous vetting specifically for "feedback-approved reviewers". Another strategy is to offer a pre-submission review, with no guarantees of scoring outcome, as a vehicle for micro-credential seekers to obtain feedback on their work before they submit, or resubmit. In traditional assessment parlance, feedback would be appropriate for formative assessment, with the goal of contributing to the immediate learning process. In contrast, summative assessments serve the purpose of making a determination about competencies or skills, and beyond the outcome are not designed to generate rich feedback to the learner. While an ecosystem that advances professional learning needs to offer opportunities for growth, the certification of competency might be more easily carried out separately from the process of gaining competency.

Another issue to consider, raised briefly earlier in this paper, refers to the reviewer recruitment and training process. There is a need to balance rigor of reviewer training with reasonable time investment and incentives for becoming a reviewer. Where micro-

credentials are organized into stacks and share characteristics of competencies, it may be possible to streamline the training process such that individuals could become certified reviewers for an entire stack of micro-credentials at once. A related consideration is the role that domain-specific expertise plays in the review process. For general competencies such as performance assessment, project-based learning, and deeper learning strategies, their implementation might be better reviewed by those with subject-specific expertise and experience, as in the case of a calculus teacher who submits a project-based learning module focused on rate of change.

Finally, research and analytics represent an exciting area for generating new insights into the micro-credential implementation process, and for documenting the ways that this approach to professional learning has positive impacts on classroom practice. As adoption of micro-credentials grows, the volume of data will enable detection of increasingly sophisticated patterns and trends regarding reliability, validity, and impact. To prepare for these opportunities, micro-credential platforms need to incorporate data architectures that anticipate the kinds of data that will be important, as well as the formats in which they are captured to facilitate analysis and reporting.