The EDUCAUSE Learning Initiative Fall Focus Session
Supporting Learning with Technology: Assessment's Role

The aim of the ELI Fall Focus Session, "Supporting Learning with Technology: Assessment's Role," was to present both the underlying principles and practical methods for establishing evidentiary frameworks for enhancing student learning. What follows are brief summaries of topics covered in each of the general sessions. Please see the event proceedings Web page <http://www.educause.edu/Proceedings/10532> for presentations, resources, and podcasts from the sessions.

How Do You Know if IT Is Helping Advance Student Learning?
Karen Swan
<http://www.educause.edu/ELI063/Program/10494?PRODUCT_CODE=ELI063/GS02>
The fundamental goal of assessing educational technology is simple enough to grasp: improving student learning and success. In her presentation, Karen Swan, research professor at Kent State, moved into the more complex realm of performing assessment and the important but less obvious concepts, processes, and nuances that are involved in the endeavor.

On a conceptual level, there has been a longstanding debate about the effects of technology on the learning process. One of the more prominent debates on this topic has been between the education scholars Richard Clark and Robert Kozma. Does technology, as Clark argues, serve simply as a conduit for information, acting like a delivery truck that does nothing more than carry food to the store? If the food the truck was carrying was ice cream, wouldn't the truck itself matter? Kozma follows this side of the argument, attesting that the kind of learning that technology supports may vary. As for the studies that have been done to gauge the differences between online learning and face-to-face learning and between computer-assisted instruction and traditional learning, these have shown us that no significant differences exist between them.

Zooming in to a more granular and practical level, however, illuminates the existence of important differences between specific kinds of learning, people, and contexts. It is assessment questions at this level that we need to focus on.

Learning outcomes are probably the most common focal point of assessment. Among the important learning outcomes to look at is student satisfaction: If the student is not satisfied, she will not be engaged, and if she is not engaged, she will not learn. Retention is also important, for if the student is not in class, he probably is not learning. Student success, usually measured in grades and proficiencies (often determined by discipline or industry group) also fall within this group. Performance is another important outcome to measure; how students can perform in the real world will affect their prospects for job placement, retention, and advancement.

A final outcome to be measured is student achievement, which, unlike student success and grades, signifies how well a student should be able to do *after* finishing a course or program. A helpful approach to defining achievement is offered by Wiggins and McTighe in their book *Understanding by Design*. They categorize the kinds of learning that can occur into three areas: *things that are worth being familiar with, things that are important to know and do*, and *enduring understanding*. The latter was characterized as the essence of a discipline or topic, something that should be understood five years after the course or program. They point out that *things that are worth being familiar with* are often used to measure student progress because they are easy to measure—for example, whether a student read a particular text. But such things really do not have lasting impact on student success. *Enduring learning* and *things that are important to know and do*, while harder to measure, are truly importance to assess.

What kinds of assessment will work for getting at enduring understanding? Generally, quizzes and tests will not. Performance tasks and projects will—but how can they be used to for assessment? One method is rubrics. Rubrics can clearly define a course's important learning outcomes and key stages in reaching them. In addition to using them to measure student learning, they also can help students themselves see what they need to learn and how well they are learning it. Also effective are e-portfolios—collections of students work that show evidence of learning in a single course or throughout a program and demonstrate progress over time. They can be linked to specific learning goals and also show student reflection about their own work. Clickers too can be used to assess deep understanding and have the additional benefits of increasing student interactivity and reducing their stress surrounding test taking.

All in all, outcomes are not the only component of learning—and perhaps not the most important—to assess. They do not indicate how or why the given results were obtained. What provides such information is assessing the inputs and processes involved.

At the front end of the learning process are learning inputs. Among those that are important to measure are learner characteristics, which include demographic information, gender, ethnicity, socioeconomic background, academic level and discipline, and technology comfort level. Another key learner characteristic input is learning style. This encompasses not only the well-known visual, auditory, and kinesthetic trio of styles but also personality characteristics identified by Myers Briggs, Kolb, and other personality type models. It is worth noting that learning styles may also differ across academic disciplines. Yet another set of learner characteristics is manifested in the reactive behavior patterns currently being researched at the University of Central Florida.

Additional inputs to asses include course design (see the survey instruments offered by the TLT Group and Quality Matters), learning resources (support materials and services available to students such as the library, writing center, or technology center), and the professional development of faculty.

The critical link between inputs and outcomes are learning processes, which are often overlooked in assessment efforts. Among the important processes to evaluate are pedagogy (the strategies and techniques instructors use), student interaction (with content, with teachers, among themselves, and with technology), and course-based assessment. Regarding the latter, it is worth noting that the choice of what is assessed can affect or change altogether the learning outcomes. Determining which assessment method effects the desired outcomes is therefore important.

Processes, because they involve so many variables, can be a challenge to assess. A study of the effects of ubiquitous technology on teaching and learning being conducted by Swan and others at Kent State illustrates one way to simplify the endeavor. In that study, the sheer number of technologies that warranted consideration made a comprehensive study overwhelming at best. The focus was narrowed to three select areas that have distinctly begun to show the effects of the ubiquitous computing environment: exploring how students learn by looking at the many kinds of external representations of learning enabled by technology, examining students' conceptualizations of knowledge through semiclinical interviews of students as they work with their external representations of knowledge, and studying student uses of knowledge—namely, the increased activity in small groups, collaborations, and other social forms of learning witnessed in technology-based learning.

Another approach to studying process is the communities of inquiry model developed by Anderson and Garrison, which looks at what students learn online, especially in online discussion, and focuses on the intersection of three kinds of presence in online interaction:

- *cognitive presence*—ideas and concepts of course

- *teaching presence*—things the teacher does to support learning
- *social presence*—community

Each presence is assigned different categories (for example, teaching presence comprises design and organization, facilitating discourse, and direct instruction) and category-specific indicators for marking them. They can be used for content analysis of discussion or to create survey questions for students.

Asking the Right Questions by Starting with Key Decisions
Deborah Keyek-Franssen
In this session, Deborah Keyek-Franssen, associate CIO for academic technology initiatives at the University of Colorado at Boulder, walked through the process of institutional instructional technology assessment and planning at her institution.

The state of Colorado requires CU Boulder to perform an information technology strategic planning cycle every four years. In the year leading up to planning, instructional technology is heavily assessed to help the planners know the ways students and faculty are using technology and to discern where to focus attention in the coming years. Contributing to the process are the university's formal IT governance bodies as well as its strategic plan for educational technology, which demands a focus on undergraduate education and the effective use of educational technology in support of the university's mission.

In the most recent planning cycle, the planning team established an assessment framework comprised of three components, supported by related criteria and guidelines: research (obtaining quality data), planning (ensuring what was collected was usable), and shared governance (securing student and faculty buy-in for decisions that came out of planning process and to make the assessment pertinent and legitimate). While the first was a matter of course, the last two were included to guarantee explicit discussing and planning for them.

For the research component, the team established three criteria for data collection: validity, reliability, and trustworthiness. They wanted to be sure to ask the right questions of an appropriate sampling of the right people, through a systematically performed process. They also involved the university's institutional research group to ensure that the data-collection process was perceived as transparent and the data collected perceived as neutral—not, for instance, branded by the CIO office or information technology services.

Regarding the planning component, "curiosity questions"—assessment questions that start out as an inquiry or as a directive from an institutional governing body—were avoided, as they are not usually connected to specific outcomes or goals, though they often provide a good starting point. To transform curiosity questions into productive questions, they employed a research technique called "backward marketing." Unlike standard research, which derives its research agenda from an open-ended question, backward market research derives its research agenda from a decision that needs to be made.

For example, if beginning by posing the curiosity question, How many faculty use WebCT? the straightforward research approach would involve researching how faculty use WebCT in their teaching and why, recognizing that the findings might not be useful. Backward marketing, however, would first determine what decisions need to be made about WebCT, such as, "We want to allocate resources effectively to reduce barriers that might prevent faculty from using educational technologies such as WebCT to their fullest," which would lead to backward marketing questions such as, "What learning management system or feature are you having problems with? Aren't using but want to?"

Ultimately, in the CU Boulder planning process, the decision the team wanted to make was how to invest the university's limited resources in educational technology to have the greatest impact on teaching and learning for the greatest number of students and faculty. When they began to define the data needed to make this decision, several embedded principles guided them. For one, they wanted to focus first on teaching and learning, and then technology as a supporting tool. They also realized they wanted to benefit the majority of people in their community: the early adopters, early majority, and the earlier individuals in the late majority. Per their strategic mission, they wanted to focus on undergraduate education. And they also understood that their current options may not be the best.

The broad strategic assessment questions that eventually emerged were:

- What would you like to use but currently do not?
- What would enable that use?
- How can technology improve teaching and learning?

The technologies they opted to focus on included learning management systems, online course components, classroom technology, communications technologies, and administrative applications.

Then came the data collection, which relied most heavily on a standard part of their four-year assessment cycle process—surveys of faculty and students. They conducted two faculty surveys that tended to give similar results, one that went to 40 percent of faculty (enough for statistical validity), and a second that went to all faculty to foster a sense of engagement. The university also administered a student survey to 200 undergraduates (just under 10 percent of the total).

The faculty were asked questions such as whether they use laptops as a replacement for the overhead projector or chalk. Sixty percent of respondents did so on a regular basis, and 75 percent thought this practice was important. Although only 33 percent used laptops for other, more innovative instructional purposes, 72 percent thought it was important to do so. This disparity between current and desired use alerted the team to areas that needed development and increased faculty attention.

The research team also asked faculty about experiences and perceptions with a specific technology, including whether they had believed it was too problematic to try, had used it but given it up because it was problematic, or simply did not perceive they needed it. They also tried to find out how faculty used online course components—such as online syllabi and readings or interactive online simulations. And they explored what could be done to encourage faculty to adopt particular technologies, such as training; access to additional time, money, or staffing; assurance ahead that it will work; or connection to tenure promotion.

Secondary but still important research outcomes included several insights for the next assessment cycle. The team will ask faculty to answer a slightly different question: What are your teaching/research challenges, and how can IT help you meet them? They also had not asked any survey questions about older technologies like slide projectors and overhead projectors. This omission gave some faculty the erroneous impression that these tools were being phased out. Faculty also expressed a desire to participate more in IT governance and decisions in the future.

The team also realized it had not planned how to weight the data elements they obtained. Qualitative factors surrounding the data collected emerged—such as the response rate for each question, the tone of responses, respondents' political position, and differences among the university's schools and colleges—and these needed to be factored in to future decision making.

Where to Look for Learning Impact
Kathy Schmidt
<http://www.educause.edu/ELI063/Program/10494?PRODUCT_CODE=ELI063/GS07>
In this session, Kathy Schmidt, director of the Faculty Learning Center at the University of Texas at Austin, offered a number of valuable examples of how technology-infused teaching practices and various assessment methods grounded in sound pedagogical research can be used to determine ways to effectively support and improve learning. The Faculty Learning Center, associated with the university's college of engineering, has used data from classroom clicker use, online class discussions, evaluations, test results, and other sources to accomplish a variety of learning improvements, from helping faculty improve their strategies for student engagement to developing a stronger student community.

When looking to gauge and improve the impact of technology on learning, some fundamental questions should be kept in mind. How do we learn? How do we learn when using technology? Who is doing what in the instructional cycle? How do you know if what you are doing works? *How People Learn*, the seminal book edited by John D. Bransford, Ann L. Brown, and Rodney R. Cocking, moves beyond speculative approaches to learning theory and applies scientific research findings about how the human brain functions to address these questions. It emphasizes effective ways to bring about learning, including working with students' preexisting understandings and using extensive examples to teach in depth. It also suggests integrating metacognitive skills into the curriculum to teach students how they learn.

The Faculty Innovation Center employs a model based on the *How People Learn* (HPL) framework that centers on four areas:
- focus on learners and tailor instruction to their needs,
- arrange instructional content around big ideas and core concepts,
- allow students to revise and improve through assessment, and
- foster intellectual camaraderie and community.

In this model, students develop ideas for solving in-depth problems, test their own solutions, and receive timely feedback from instructors. Instructors are encouraged to clearly define the structure of the class and the outcomes students will be working toward.

The UT engineering school offers several practical illustrations of the HPL model. One example is the bioengineering program, which seeks to instill ethics and professional responsibility in its students through the concept of "adaptive expertise"—applying learning and experience to unanticipated situations. UT adopted a pedagogical approach called STAR.Legacy Cycle (STAR stands for Software Technology for Action and Reflection) to effect this kind of learning by encouraging active engagement and independent thought. In essence, this method presents students with a problem and asks them to develop their own ideas about how to solve it. Students in this program discussed the problem as a class, fine-tuned their solutions through independent research, tested their solutions, and ultimately went public with their final results.

To discover the impact of this approach, the school conducted a study that compared a traditional lecture model with the STAR.Legacy/HPL model and found that the latter did a better job of developing adaptive expertise. While both sets of students learned the courses' factual information and could develop arguments on individual topics, the HPL students also were able to integrate multiple perspectives into a broader ethical argument.

Another example focuses on the cultivation of community in a lecture-based instructional setting. The Faculty Learning Center was approached by a professor who believed that class time should be dedicated to lectures and the "intellectual heavy lifting" that comes with that format, but who also wanted to act on research showing that community was important to learning. Because only so much class time was available for presenting important concepts, the professor did not know how to fit in community-building activities. In response, the center developed an e-portfolio system for the class that gave student groups a way to review and give multiple rounds of feedback to each other online, outside class. The professor ultimately was able to determine from the feedback whether students were engaged, learning, and creating community. The center also helped the professor develop a student self-assessment survey to determine how they were responding to his classes overall, which showed positive results.

A third example involves a statistics class that had been unpopular with students, particularly because the content seemed irrelevant to the real world. The class's professor was keen on improving the situation and introducing more collaborative and computing activities into the course. To determine the steps that should be taken, the Faculty Innovation Center staff studied current classroom technology (PowerPoint, Blackboard, clickers), observed the classes and labs, facilitated self-report surveys, reviewed online class resource access statistics, and used Bloom's taxonomy. The professor has used the findings to improve her class.

Assessment Methodologies
Patsy Moskal and Joni Spurlin
<http://www.educause.edu/ELI063/Program/10494?PRODUCT_CODE=ELI063/GS08>
Making assessment manageable was the underlying theme of this very practical session on assessment methods with Patsy Moskal, associate director of the Research Initiative for Teaching Effectiveness at the University of Central Florida, and Joni Spurlin, university director of assessment at North Carolina State University.

Starting with the basics, what should be assessed? First, one must determine which level of learning to be assessed. At the student level, assessment gauges how individual students are performing. On the course level, it measures how well students are meeting course objectives. Program-level assessment looks at how well students integrate knowledge from courses into a fuller professional understanding. In addition, assessment of educational technology should look at learning outcomes rather than particular technology attributes.

The assessment cycle usually follows these steps:

1. Define research questions
2. Determine the approach to use for finding answers
3. Proceed with the research and data gathering
4. Analyze the collected data and interpret the results
5. Make decisions based on these findings

Although it may be common sense, it is worth emphasizing the value not only of planning ahead for assessment but also of integrating them into teaching and learning projects. It is also important to consider the assessment research questions thoroughly in advance so that data that is collected from them can be used. Uncollected data cannot be analyzed, and going back to ask additional questions can tax respondents' already short time. Also, useful data may have already been collected elsewhere that may assist planning assessment research.

Working as a team rather than alone and seeking advice from experts on campus can lead to better results. Among the people who may be valuable to involve or consult are assessment

professionals, IT professionals, instructional designers, decision makers, faculty with assessment or evaluation background (in sociology, education, and so forth), the institutional research/planning office, and the faculty development center.

A potentially daunting element of assessment is the need to go before one's institutional review board to gain research approval. It may be helpful to seek advice from colleagues who have been through the process before. Also potentially intimidating is navigating the complexities and demands of statistical research methods. Sometimes, however, data can only be gathered, and statistical significance is not required or possible.

Once data have been collected and analyzed, they will need to be interpreted and presented in way that will be easily understood by those who hear it, be they faculty or administrators. Visuals can be helpful in this process.

Establishing assessment as an ongoing process is important. Technology is not a static thing, nor is the campus or its occupants. Also, regularly obtaining assessment feedback will enable continuous improvements.

There are two fundamental approaches to assessment: indirect, which focuses on learners' opinions of their knowledge, skills, and qualities, and direct, which focuses on the manifestations of these aspects of learning. Among the most commonly used indirect methods are surveys, which are relatively easy and inexpensive to administer. Unlike other indirect methods such as focus groups or interviews, they offer the researcher firm control over the direction of inquiry and, because the same set of questions is asked of the entire research group, research bias is less of an issue.

Surveys administered in person tend to garner stronger response than those administered electronically, largely because survey respondents are able to directly see that the researchers care about their answers. Also, including open-ended questions may add valuable depth to response data. While the answers may be more difficult to analyze, details will be revealed in them that would be impossible to obtain in the survey format otherwise.

How the survey questions are worded will strongly influence participants' responses . Care must be taken to phrase questions in a way that elicits responses that are as honest as possible and that address exactly what one wants to know. They should not contain wording that implies that respondents could negatively be affected by the answers or results, such as an impact on grades. They should also avoid wording that is widely open to interpretation. For example, "very good" may mean superlatively good to some people and simply slightly better than good to others.

Focus groups, in which a moderator interviews a group of selected individuals, are another way to conduct indirect assessment research. While they may not be as time efficient as surveys, they are more so than one-on-one interviews. They also can generate a greater volume of information and sometimes deeper information than either of those methods, because of the open-ended format and the fact that group participants may draw out information from each other in a way that an interviewer alone could not. It is especially important to employ an effective and personable moderator to keep the discussion flowing and focused on the research topic, as well as to put the group members at ease. A downside of this method is that the large amount of resulting data can be overwhelming.

Among the various direct assessment methods available are rubrics, which are a printed set of scoring guidelines for evaluating work. They generally require selecting an outcome to assess, define, and "operationalize" it (noting all the skills, qualities, and knowledge that go into it), set

different developmental stages or "dimensions" for each operational step, plot all the skills and their dimensions in a grid, and ultimately use them to score performance. Examples of various rubrics can be found easily by doing a simple online search.

Rubrics can be effective for assessing how different instructional approaches affect a particular learning outcome. For example, they can help measure the impacts of differing pedagogical approaches on student critical thinking. The first step in developing such a rubric is operationalizing the critical thinking elements, such as analyzing facts, gathering and evaluating evidence, and transferring thinking processes from one subject to another, followed by defining roughly four dimensions in the development of each element, which would be plotted in a grid and used for assessment.

Along with developing the rubric itself, faculty must be prepared to use it. In the example above, this would require working with faculty members offering courses whose critical thinking outcomes were going to be assessed. To create measurable standards in the rubric, participating faculty would need to assign similar projects involving critical thinking in each of their courses. They would incorporate into these projects the use of various instructional technologies, such as discussion boards or simulations. Then they would use the rubric to assess all their students' projects, and the results would be tallied, averaged, and analyzed according to the technologies used.

Strengths of the rubric method are that it can be used by others besides the instructor and can be reused when assessing other technologies. Rubrics also very clearly define what is to be measured and support assessment of higher cognitive abilities, but they will be only as good as the definitions and dimensions defined for them—and developing them can take considerable time and experience. The process also can initially be somewhat demanding of faculty because of the course adjustments they may need to make.

Another direct assessment approach is course-based or embedded assessment, which involves collecting information generated in class work about students' progress toward established learning outcomes and analyzing it with a matrix. This method can be useful for looking at how a specific technology affects different learning outcomes and for linking program outcomes to course outcomes.

Like rubrics, course-based assessment involves first operationalizing the learning outcomes to be studied, then determining how the course-based information will be collected and by whom. Once the data is collected, they are tabulated in a matrix set up with the established learning outcomes. The matrix typically will contain courses to be assessed (normally at least three) in the left column the, the learning outcomes to be measured for all the courses in the header row, and the specific assessment criteria for each course/outcome pairing in the column cells.

An example would be a scenario in which students across four classes were supposed to be able to analyze data, show engagement with their learning, and visualize the material using a particular technology. The faculty in the four classes would need to agree to a common way to measure these outcomes, such as including a set of five related questions in their final exams. The student performance on these questions would then be tabulated in a matrix set up with the three outcomes and evaluated.

This method makes efficient use of existing student work to accomplish assessment. Also, student motivation will likely be higher than it is when the assessment project exists solely for assessment, independent of grades, as they are for rubrics. In addition, the results should indicate where improvements need to be made and can be monitored over time to judge the

impact of improvements. This method also clearly connects assessment to outcomes and learning to curriculum.

As with other assessment methods, course-based assessment has downsides. The time and cost in the initial planning and design can require significant resources. It may also require extra work for faculty, who may have to tweak their course syllabi to develop common course outcomes. Also, faculty may not be willing to share their results with other participating faculty or may find it difficult to agree on methodology.

Regardless of the assessment method used, the use of educational technology creates additional outcomes that are discretely different from those that occur in traditional learning environments. It is therefore important to measure these new outcomes separately from existing outcomes, as they are not the same.

Making Evaluation and Assessment Part of the Cultures
Charles Dziuban et al.
<http://www.educause.edu/ELI063/Program/10494?PRODUCT_CODE=ELI063/GS09>
One of the motifs of the focus session was the importance of making assessment an ongoing practice. This session with Chuck Dziuban, director of the Research Initiative for Teaching Effectiveness at the University of Central Florida, and others touched on ways to work assessment into the institutional culture so that it is expected, facilitated, and used.

Although Walt Disney World may have radically different goals from higher education's, the company's use of data in daily decision making was highlighted as a valuable exemplar for integrating assessment into daily institutional operations. Every evening, the company analyzes its occupancy data using a Boolean algorithm to determine for how long each room needs to be sold to make the maximum profit. This information is provided each morning to the hotels' staff, who then use it to determine which rooms to sell. In this way, data are inseparable from daily decision making.

"Facilitative leadership" is a leadership model that emphasizes power sharing—fostering collective engagement in problem solving and improvement across the organization. Several of its principles are applicable to the development of an assessment culture. Involving as many people as possible and appropriate in the assessment process and cultivating a shared inspiration for fulfilling the assessment vision will help build interest and support in the process. Clearing the way for those involved and making it easier for them to integrate assessment into their routine will also help. In addition, those who are involved in the assessment cycle will feel more involved and engaged if the emphasis is not only on process and results but also on relationships. Along the same lines, by bringing out the best in others and behaving in a way that models collaboration and respect, others will be more willing to participate. Finally, celebrating successful assessment accomplishments will be appreciated and build support.

Don Carter, director of e-learning at Northern Arizona University, shared a description of the recent growth of assessment at his institution. About five years ago, state grant monies were used to create the Center for Research, Assessment, and Development of Learning in Electronic Environments, known more commonly as the E-Learning Initiative. The center, which includes an Office of Academic Assessment, has a faculty fellowship program that supports research on pedagogies and practice, implementing technologies, addressing the needs of learners, and exploring emerging and collaborative learning tools. The center also offers 10 annual minigrants to faculty to focus on ways to improve classes that have high enrollment, higher DFW rates, or other similar characteristics; each grant project includes an assessment plan. Another area in which the center has been involved is in consolidating efforts with other related university sectors

to create a teaching and learning support services Web site. This site offers links to assessment, distance learning, e-learning, faculty development, ITS, library, and liberal studies programs. And behind the scenes, these groups work collaboratively and speak collectively with faculty. Beyond the center, a longer-standing faculty-controlled University Assessment Committee seeks to ensure that assessment plans are integrated into the university's degree programs.

Another example of thriving assessment culture was spotlighted by Linda Jorn, director of the Digital Media Center (part of the Office of Information Technology) at the University of Minnesota System. Although she began in a half-time position in 1995, the center now employs 17 consultants to provide customized support to faculty as well as three researchers who focus solely on evaluation. The mission of the unit is to use technology well; partner with the school's colleges and vice provost office; and focus on evaluation and research in designing technology for learning environments.

Back in 2002, the center focused on administering institutional assessment and faculty technology surveys that examined faculty needs and on improving customer support. Through their findings, center staff learned that faculty feared losing interactivity in the online environment and wanted to learn more about online pedagogy. On the other hand, students felt faculty were not using enough technology and were not using it well. This led the team to shift from an emphasis on faculty workshops to faculty development programs. They now offer seven such programs for faculty. Assessment is embedded within them, and collaboration is emphasized throughout. The scholarship of teaching is stressed as well, and faculty are encouraged to put it into practice.

One of the faculty development programs the center offered was developed in response to faculty interest in improving large-enrollment courses. Linda's unit set up teams for 10 large-enrollment courses; each was comprised of a faculty member, graduate student, undergraduate student, and consultant. These teams, usually together for three years, met to set and implement strategies, and reflect on and revise them over the course of each semester. All the teams met together monthly to review and discuss their work, get feedback, and seek guidance from the consultants.

Echoing other presenters' sentiments about the importance of establishing good metaphors for assessment to communicate its essential aims, Gardner Campbell, professor of English at the University of Mary Washington, offered closing thoughts about the metaphors we associate with assessment. He suggested the word "witness," whose various meanings include wisdom or knowledge; an accountable careful observer; and to stand as an observer or to observe. He spoke of the quietly moving example of a colleague at the university who had recently passed away. Tadesse Adera, an English professor at the university, had had the appearance of a professor whose teaching methods and accomplishments were unexceptional. Yet on his passing, his students bore passionate testimony to the effectiveness of his teaching, including one who sent a letter sharing how much Adera had contributed to his life. In this way, what could not be measured could be witnessed and attested to.