Distilling Competency Information

White Paper – Draft 4 March 2005

Claude Ostyn
Ostyn Consulting

Abstract

Automated competency tracking and management in the context of performance support, training and adaptive online learning requires a systematic way to track competency data for individuals and teams. However, this competency data may come from a variety of sources and in many different formats. This paper proposes a simple model to distill the data into simple standard formats that can then be used as input to guide decisions, skill gap analysis, and adaptive learning. The model addresses a variety of enterprise policies, as well as credibility issues, sanity measures to avoid corruption by unreliable data, auditing, and general security and privacy requirements. The main goal, however, is to enable this functionality with the simplest possible model while taking advantage of standards that already exist. This white paper may result in specific proposals for new standards or for inclusion in existing standards projects.

Contents

Background

Defining competencies

Competencies are a very vast subject complicated by very strong opinions and cultural traditions. High level competencies are highly personal and essentially impossible to define and describe—the sum is greater than the parts. For example, the competency of a charismatic leader can be analyzed, but such an analysis will never capture some intangible aspects of what makes this person capable of such performance. However there is a lot about competency information and competency data that can be described or captured to enable systematic performance improvement, performance support and training. Standard data models and clearly defined processes or workflows can then enable various forms of automation, such as computer storage of personal competency profiles, skill gap analysis, individualized performance support and training, and decision support for individuals and their managers. For example, such data can be used to find learning resources designed to build the specific competencies required for a task.

The standards-based competency information described in this white paper is not intended to replace human judgment and intuition. Rather, it is intended to allow humans to make more informed decisions by leveraging existing data and knowledge, and to facilitate efficient learning workflows leading to better performance. This information can also be used to support performance directly; for example, if the competency requirements are defined in the workflow for any real world task, the information models proposed here support the identification of individuals or teams qualified for the task. They also support the development of the competencies required for the various aspects of the task, as well as just in time performance support if should become necessary.

Reusable Competency Definitions (RCD)

A reusable competency definition captures "the part of competency information that may be reused for more than one person in one or more contexts and possibly with different metrics"

- The definition may define as skill, knowledge, attitude, learning objective.
- The definition may be for a facet of competency (e.g. affective, psycho-motor, cognitive facets) or for a component competency of a larger competency (e.g. using the steering wheel as part of being able to drive).
- Two parts: The identifier that makes it globally unique, and the "body": Title, description, definition(s) according to one or more community-specific models, intrinsic metadata.
- A lot of existing competency definitions can be captured into a standard data format. New ones may be created directly using the format. Metadata provide a means to identify the source of the definition, as well as possible known fixed relationship with other definitions or contexts.
- IMS specification, on track to become IEEE standard, under consideration as ISO standard
- See http://ltsc.ieee.org/wg20 or IMS RDCEO specification.

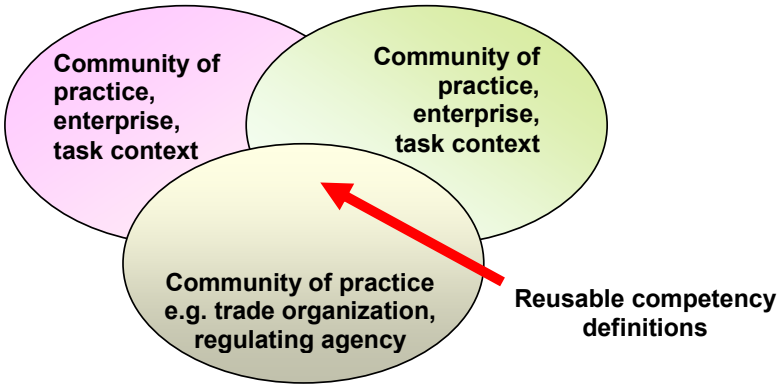

Figure 1 - Reusable competency definitions don't have to be reinvented for each context

Note that the term "reusable" does not imply that every RCD is reusable; rather, it implies that the exact same data container format can be used for all kinds of competency definitions, regardless of their scope of reusability. For example, a particular skill definition may be extremely specific to a particular piece of equipment, with very specific performance criteria. But it can be captured and referenced using the same data model as a more generic definition, such as "able to resolve conflicts rather than escalate them".

Reusable Competency Maps

A competency map is a map of RCDs. It is not made of RCDs, but rather it shows how RCDs are related. It is a structure collection of nodes that reference RCDs.

- The proposed competency map model is a taxonomy of nodes that refer RCDs. The child nodes of any node can represent facets or component competencies, or both.
- Ontologies vs. taxonomies. Small ontologies look easy, but they don't scale well. Useful ontologies are hard to build and brittle. Massive taxonomies are hard and brittle. Perfect taxonomies will never exist. Small taxonomies are easy. Keep it simple.
- The same RCD may be referenced by more than one competency map, and/or by more than one node in a competency map (see Figure 2).
- Different communities of practice may map the components and/or facets of a competency in different ways. This is a reality that cannot be overcome. Embrace it, don't fight it.
- Competency map provides part of the context for a RCD. For example, it can help assess a competency in the context of related competencies.
- One or more competency maps may be specifically referenced in metadata for a RCD, using a classification element. This may even specify the "taxonpath" that specifies the path to the node that references this RCD in the taxonomy.

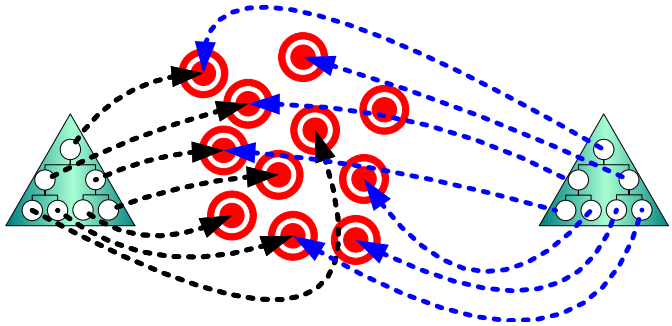

Figure 2 – Different maps reusing some common competency definitions

Task models

Task models come in many different flavors, according to many different models. Some are hierarchical, others are ontologies. Some define a qualification; some are specific to a domain, e.g. "Radio equipment operation". Some are very specific, e.g. "repair automobile automatic transmission model XYZ", others can be very vast and describe an occupation, e.g. "Accountant" Some are defined as workflows; others, like the US Army MOS, are defined as hierarchical layers of disparate elements.

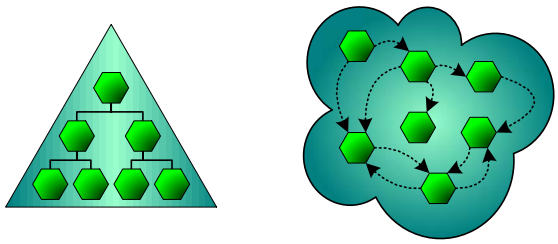

Figure 3 - Task models can be simple hierarchies of tasks, or they can be very complex models

It does not appear possible to standardize task models across communities of practice. However, the existence of more or less explicit task models is assumed in this paper. In the sample

scenarios that were examined to validate the concepts in this paper, consideration was always given to how standard data and processes can be applied to real world task or workflows.

Evidence of competency

Evidence of competency is usually based on some form of assessment. Assessments also come in many different flavors. For the purpose of this paper, we will use a very general definition of assessment:

> ***An assessment is any form of judgment, automated or not, that produces results that that can be captured in a digital data record.***

For example, if a reviewer inspects a claim of competency in a person's resume and decides that it looks credible, this is a form of assessment for which a result may be captured in an evidence record. Assessments may be performed by humans or automated. They can be simple or multi-faceted. This paper proposes a common model to distill assessment results into evidence records, regardless of the instrument or form of assessment.

Obviously, different forms of assessments merit different levels of confidence. The results of a properly conducted 360 degree assessment are clearly more credible than those from a SCORM learning object that was used on an unsecured computer somewhere on the Internet.

The model proposed in this paper assumes that there can be at least one form of assessment for every defined competency. In practice, assessments often evaluate multiple competencies, or evaluate composite competencies. We assume that each of those competencies, component competencies or facets of competency can be represented by a reusable competency definition, possibly with some additional context-specific metrics. For example, an assessment may be designed or specified to evaluate whether a racing pit team is capable of changing a race car tire in less than 30 seconds. Another assessment may be required to assess whether a person achieves a set level of proficiency.

Distilling competency evidence

The distillation process for a single competency or competency facet is summarized in Figure 1. This distillation process is used for each reusable competency definition for which one wants create a corresponding competency record.

1. Some form of assessment of a person or team's proficiency in a specific competency or competency facet produces results. Different forms of assessment produce different kinds of results. There is no single standard for assessment results.
2. The assessment results are distilled into standard evidence records. Results may come in at different times from different sources.
3. The most credible evidence record is used to create or update a competency record that states whether an individual or team satisfies the requirements for the competency, and at which proficiency level.

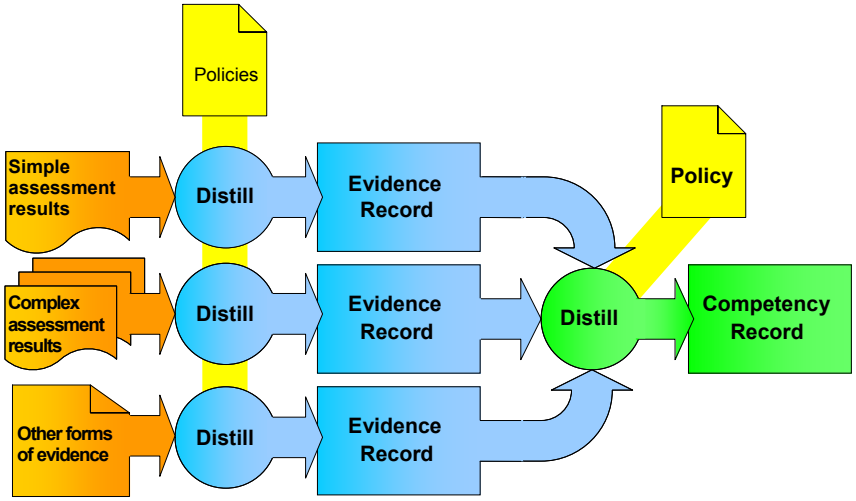

Figure 4 - Summary of the competency evidence distillation process

Standard evidence records and a competency record can also be qualified with a confidence rating. This confidence rating is determined by policies. The proposed standards do not specify what those policies are, since they will differ with enterprises and communities of practice. However the proposed model provides a way to capture the result of applying those policies, and to update the data when policies change.

Competency records

A competency record states that:

- A person (or a team)
- Has a known competency status, which may be qualified with a proficiency level, regarding a particular competency or facet defined in a reusable competency definition
- This determination is based on one or more records of evidence, which can be looked up
- Based on enterprise policy, the evidence is considered more or less credible

Depending on enterprise policy, the competency record is updated when new evidence records are considered, or when the confidence accorded to one or more forms of evidence changes.

For example, one enterprise's policy might state that competency records are updated on an annual basis, based on the evidence collected during the year. Another enterprise policy might state that the competency records should be updated on an ongoing basis, as soon as new evidence becomes available.

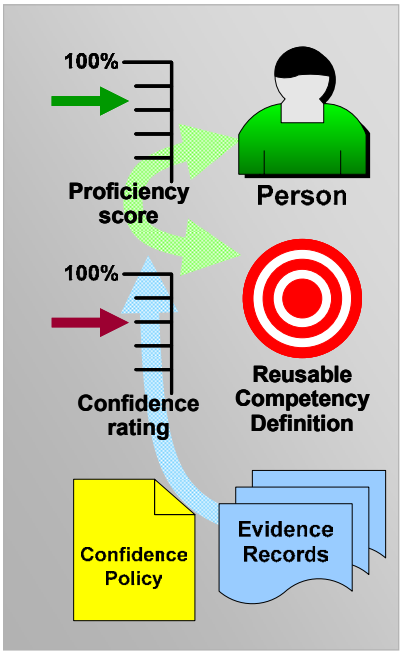

Figure 5 - Competency record

Competency evidence

A competency record must be supported by evidence. There may be one or many evidence records for the same competency. All these records reference the same competency definition and all the evidence records have the same structure, regardless of the actual source of evidence.

A competency evidence record states that:

- A person (or team) has a known proficiency in a particular competency defined in a reusable competency definition
- The evidence that supports this determination is available in the form of assessment results that can be looked up
- The results of a particular assessment constitute the evidence, and have been distilled into a proficiency status and score.

Based on enterprise policy, different confidence ratings may be associated with different sources of evidence and types of evidence. Typically, the source will be some form of assessment, but a distillation policy may also allow direct creation of an evidence record through an "on the fly" assessment of some other kind of source.

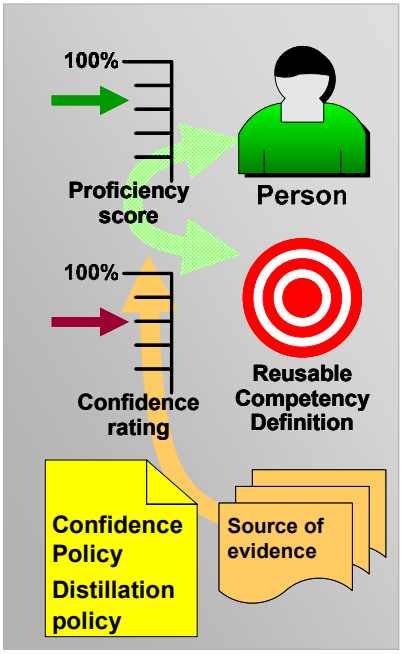

Figure 6 - Evidence record

Source of competency evidence

A source of competency evidence

- Is a tangible document, digital or not, that can be used as the source for an evidence record.
- Can be referenced in a persistent way in an evidence record (e.g. through some kind of locator or URI)
- Can take many forms, e.g. a certificate, a transcript, a driver's license, a note inserted in a personnel record, a set of SCORM runtime tracking data according to IEEE 1484.11.1, a portfolio, etc.
- May consist of, or contain formal assessment results, but not necessarily
- Assessment results may conform to a standard, or be easily distilled to a standard format.
- Is typically retained in some secure storage to allow auditing.
- May have an expiration or destruction date.
- Is more or less trustworthy and credible.

Figure 7 – Sources of evidence

There is no point in trying to standardize sources of evidence. However, a standard vocabulary for types of sources of evidence would be useful.

How a source of evidence is distilled into one or more evidence records depends on the type of source. For example, various aspects of a portfolio might be assessed for signs of competencies defined in various competency definitions. Depending on policy requirements, the result of this may be captured in formal assessment records (see below) that can then be distilled into evidence records with a full audit trail, or directly summarized in evidence records that point back to the portfolio as the source, with a confidence rating that depends on who is creating the evidence record.

Assessment result

An assessment result

- Is a tangible document that can be used as the source for one or more evidence records
- References one or more reusable competency definitions, or can be mapped easily to one or more reusable competency definitions.
- Results from one of may forms or assessment, that may be formal or informal
- Is typically retained in some secure storage to allow auditing.
- May have an expiration or destruction date.
- Is more or less trustworthy and credible.
- There is no point in trying to standardize assessment results since there are many kinds of assessments, and new ones are invented every day.

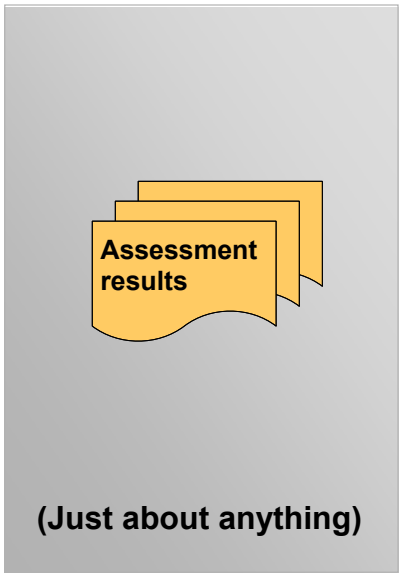

Figure 8 – Assessment results

Different assessment instruments may be treated differently as sources of evidence. For example, the results reported by an online tutorial are typically less credible than those reported by a supervisor who can observe a person's performance of the task supported by the competency. Given the same assessment instrument, results reported by an unknown third party are less credible than those reported by an enterprise insider, and so on.

Enforcing sanity policies

The quality of assessment data is of course very important. This is why it is important to be able to assign confidence rating to any data that purports to state that someone is competent in anything. Policies must guide the determination of confidence ratings in the distillation of competency evidence. Different sources of evidence are more or less credible. A self-report in a resume is less credible than a validated, proctored exam result. An assessment by a manager is less credible than a well conducted 360 degree assessment, because the manager may have a grudge or an affective attachment to the person being assessed.

Generally speaking, one should be able to guard the content of a competency record against corruption by unreliable data. Even so, the competency record itself should be rated by how much it can be trusted. For example, one can create a competency record when the only source of evidence is a college degree, but this cannot be trusted as well as when there is evidence of actual performance.

Simple to determine the confidence ratings in the distillation process will be examined in more detail in a subsequent section of this paper.

Sanity policy scenarios

Scenarios: Distilling results from online assessments

In a training and knowledge management context, it is very important to be able to access competency information. For example, SCORM allows the creation of adaptive learning activity packages that can skip the competencies already mastered by the learner.

On the other hand, one must exercise care when automated processes are allowed to update competency data. For example, SCORM can report status and proficiency levels for specific competencies. One would not allow the SCORM data to modify a person's competency record without at least some filtering through policy. It is not wise to allow the records that govern a person's career or promotion opportunities to be tainted by an unproctored, online test that may or may not have been validated. Figure 9 illustrates how competency data might flow in and out of a SCORM delivery environment.

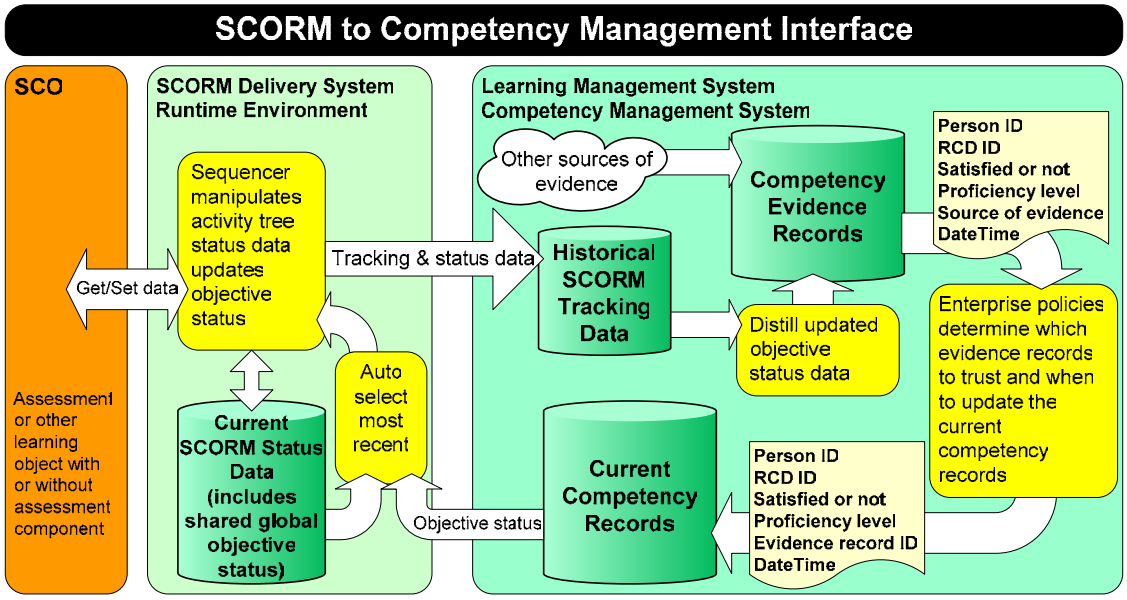

Figure 9 - Enforcing sanity in handling competency evidence from SCORM

Scenario: Distilling results from a recruitment interview

A formal recruitment interview typically will result in various assessment notes. In the more formal cases, the interview is structured to assess specific competencies or facets of competency according to a preset outline. In less formal cases, there is still usually an attempt to address key issues relevant to specific competencies of facets of competencies. The result of the interview may be captured in various formats, ranging from scribbled notes to formal data entry in a computer-based form. An interview often follows some assessment of the applicant's resume or portfolio. That pre-assessment may have been used for screening only, or to raise flags for interview questions, and the results of the pre-assessment may be recorded along with the results from the interview and recruitment process.

After an interview there may be one or more sources of evidence for evidence records:

- Interview notes or other formal results can be matched to specific competency definitions to distill evidence records with a confidence rating that may vary depending on the seniority or qualifications of the interviewer.
- Resume statements can be matched to specific competency definitions to distill evidence records. These will typically have a low confidence rating, since the applicant's own claims should always be considered with caution.
- The applicant portfolio can be assessed by an expert for signs of mastery of various competencies. The result of this assessment can be a collection of evidence records with a confidence rating that depends on the seniority or qualifications of the assessor.

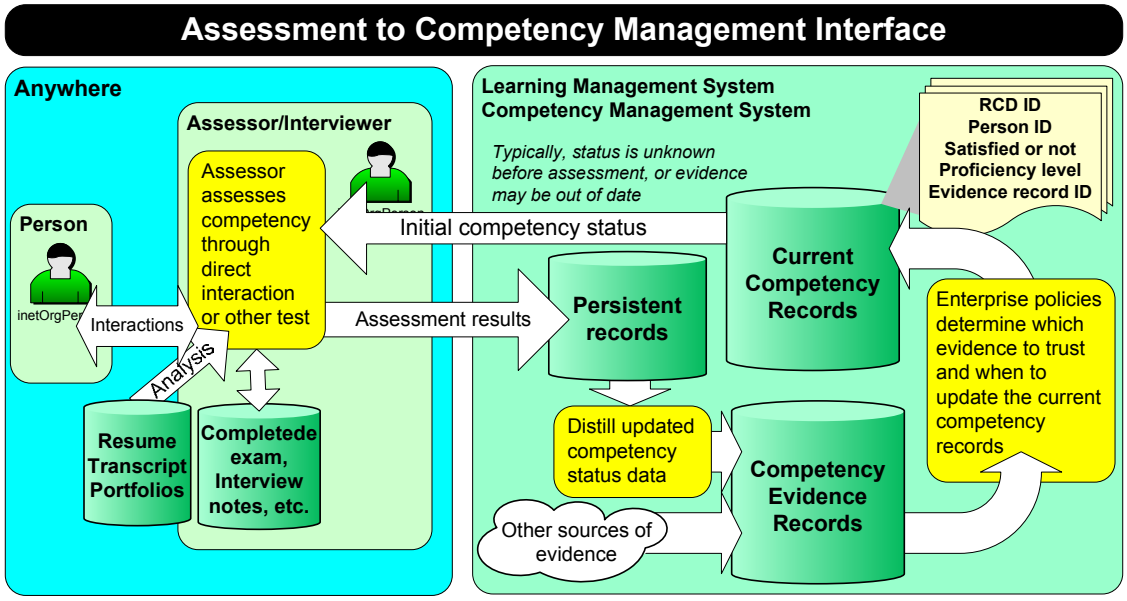

Figure 10 – Distilling competency information from the interview process

At the end of the day, if there is more than one evidence record for any competency, the record with the highest confidence rating is used to distill a competency record. The competency record will carry a confidence rating that can then be used to decide whether it can be trusted for a hiring decision. For example, in an attempt to automatically roll up the competency record data according to a competency map that represents the competency requirements for a particular job, an automated system may detect that a particular competency is not supported by sufficiently credible evidence. This in turn may trigger a follow-up activity to fill this information gap in a way that generates more credible evidence. This might be as simple as a phone call from a trusted personnel specialist to the applicant to ask a question that was missed in the interview, or verification of a claimed qualification with a university or previous employer.

Scenario: Distilling evidence from multiple sources

When data from multiple sources results in more than one evidence record for any competency definition, the record with the highest confidence rating is used to create or update a single competency record. Only evidence records that have not expired are considered, of course. If two evidence records have the same confidence rating, the most recent one is used.

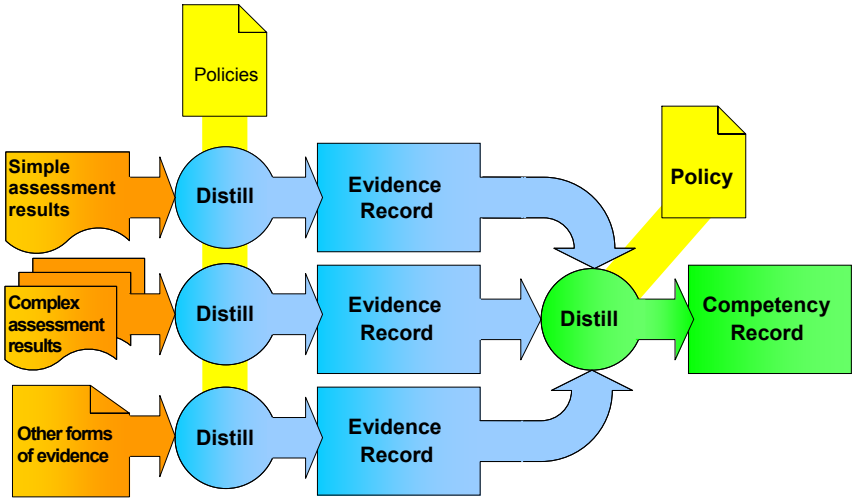

Figure 11 - Summary of the competency evidence distillation process

Scenario: Distilling expiration dates

Some evidence records may include an expiration date or other time limit. This can be carried over into the competency record. When the competency record expires, the existing evidence records can be searched for the record with the highest confidence rating that has not expired yet. If two evidence records have the same confidence rating, the most recent one is used.

An enterprise policy may dictate that the competency data should be updated as soon as new evidence is available. In that case, when a new evidence record is created, or an existing evidence record is updated, this can trigger the distillation into a competency record. If the confidence rating for the new evidence is higher or equal to the confidence rating of the existing competency record, the competency record is updated to use the new evidence. Otherwise the new evidence has no effect.

Assessment instruments

Assessment instruments come into many flavors, more or less formal, and more or less complex. The validity of assessment instruments is extremely variable. Some assessment instruments go through very thorough design processes and validity testing using longitudinal studies to verify whether the assessment results are reliable predictors of future performances. Other assessment instruments are designed on paper napkins in a restaurant; they might consist of a list of questions to ask in an interview. Sometimes assessment instruments are checks on learning following a presentation in a computer-based training module. Assessments may be administered in very secure, trustworthy conditions. Others may be administered online with no guarantee that the assess person is actually sitting in front of the computer.

Roughly speaking, one could classify assessments along three dimensions:

- Complexity – is this assessment for a single competency, or for a more structured collection of competencies and/or competency facets?
- Results format – at one end of the scale is whether this assessment produces results that can be captured using a standard data model, such as SCORM interactions. At the other end of the scale would be a result format that is very specific to the instrument, with no hope of mapping to a standard format
- Trustworthiness – the extremes are whether the assessment trustworthy, i.e. a formally validated assessment instrument or some assessment by a qualified and highly trusted person, or whether the results of the assessment should be treated with caution as informative data with limited reliability.

Data standards from learning technology

The learning technology community has come up with useful standard data models that allow the capture and exchange of assessment results. These results can then be used directly in an automated distillation process to produce standardized evidence records. These standards are:

- SCORM 2004 "cmi" data model. This data model is an implementation of the IEEE 1484.11.1 standard (see below). It is used in SCORM for communication between Shareable Content Objects (SCOs) and a runtime environment. SCOs can teach, assess or both. The same SCO may assess multiple competencies by referencing RCDs as objectives.
- 1484.11.1-2004 IEEE Standard for Learning Technology--Data Model for Content to Learning Management System Communication. This standard describes a data model to support the interchange of data elements and their values between a content object and a runtime service (RTS). It is based on a current industry practice called “computer managed instruction” (CMI). The work on which this Standard is based was developed to support a client/server environment in which a learning technology system, generically called a learning management system (LMS), delivers digital content, called content objects, to learners. The data model supports learner data and preferences, interactions, objectives, content-object entry, exit, and status information, time parameters, and scores.
- SCORM 2004 packages. The Sequencing and Navigation specification in SCORM 2004 may update information about specific objectives by rolling up results from the learning activities defined in the package. The learning activities do use SCOs to collect the data. A SCORM 2004 package is not limited to learning; it can be used to specify and deliver a simple or structured assessment in any SCORM 2004 conformant delivery environment. The assessment can be adaptive and use a variety of interaction techniques to elicit useful responses.
- IMS QTI 2.0 specification. See http://imsglobal.org for details.

As discussed above, sanity policies should apply when using standardized assessment data, depending on the source. Note that the IEEE 1484.11.1 standard was designed so that it could be used to transfer and archive assessment data in a variety of settings, not just SCORM content delivery. For example, you will find that the responses to a many paper-based or interview-based assessment instruments, as well as the results of the evaluation of those responses, can be captured in the IEEE 1484.11.1 data model.

Other data standards

There is a long tradition of development and use of assessment instruments generating more or less standardized results. Every community of practice that has such results should be able to define a method to distill those results into standard evidence records.

Assessment from existing data

Existing data in HR databases, learning management systems, and other records can be used as source of evidence. Depending on the format, an assessment instrument can be designed to capture this data. In this case the assessment does not involve the individual or team being assessed, but rather it is an assessment of the existing evidence to massage it into a form that can be used to distill evidence and provide an audit record of how this was done. Such assessment instruments might range from simple forms that are used to generate a persistent record documenting the source of data, and that require transcription by a human operator; to fully automated processes that might use web service interface to extract and massage data from existing databases. We are calling this an assessment instrument because in many case some form of judgment or evaluation may be necessary to determine whether the data is reliable and relevant.

In all cases, what matters here is that there is a way to map existing data to reusable competency definitions so that standard evidence records can be distilled from these various forms of assessment. It may necessary to create the reusable competency definitions based on the available data. For example, the facets of competency evaluated in a 360 degree assessment for people in a particular position in your enterprise might each be represented by a separate reusable competency definition.

Enforcing confidence policies

As described above, standard data models for competency evidence and competency record can be embodied in simple data structures. These data structures should include a confidence rating on a standardized scale. How the confidence rating is determined, however, is a matter of policy. It does not seem practical to standardize the policies, but it is practical to standardize the result of the policies. This takes the form of confidence ratings which enable the automated filtering and distillation of evidence.

Although it does not seem possible to define a general standard for policies, a policy that specifies the confidence rating to apply for various sources of evidence might look like the table shown in Figure 12.

Evidence source type	Rating	Comments
Resume statement, unverified	0.1	Self rating is better than nothing.
Online assessment, no authentication	0.2	Anyone might stand in for the person or team to be assessed.
Online assessment, identity verified	0.5	Proctoring ensured that there was no stand-in.
Peer	0.5	
Online assessment with certified validity, identity verified	0.7	High quality, scientifically validated assessment.
Resume statement, verified by HR	0.7	
Supervisor	0.7	Supervisor may have personal issues clouding judgment
Training instructor	0.7	Performance in training situation may not be entirely reliable
360 degree, unaudited	0.75	
360 degree, audited	0.9	
HR Director	0.95	
Executive	1.0	Trumps everything else

Figure 12 - Table summarizing a sample confidence policy

Competency record that results from such a distillation also includes a confidence rating. This allows automation of various decision processes, according to policy. For example, a policy might state that no promotion decision can be made based on competency records where the confidence rating is less than 70%. In some other cases, as when it is necessary to assemble a team to deal with an urgent situation, a manager might decide to use a confidence rating as low as 10% as the threshold for a first cut to identify possible team members.

Complex assessments

Complex assessments for multiple competencies, competency facets and/or component competencies can be specified in a standard way.

An assessment request specifies:

- The identifier of a reusable competency definition
- Optionally, the proficiency level at which to assess
- Optionally, a reference to a competency map that specifies components and/or facets of the competency to assess

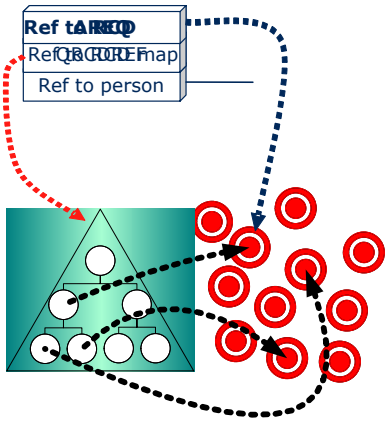

Figure 13 - Assessment request

If no person or group to assess is specified, then this assessment request can be reused over and over as needed.

In addition, an assessment request may specify one or more acceptable assessment instruments, and might specify a context in which the instrument must be used.

A target confidence rating might also be specified; but in the end it will always be the recipient of the assessment results who will decide what rating to apply when distilling the assessment results into evidence records.

An assessment may return a single result, based on the roll-up of the competencies or facets evaluated in the assessment. It might also return detailed results, with a status and possibly a score for each of the competencies or facets. A competency evidence record can be distilled from each of those results, as long as the result includes a reference to a reusable competency definition. For example, an assessment for the ability to do a task may include knowledge and skill sub-competencies. An individual might have the knowledge but lack the skill, in which case the roll-up would indicate a lack of competency for the task, but the detailed evidence would indicate that only the skill needs training or remediation.

Reusing or defining competency definitions

Enterprise policies and workflows determine whether and how reusable competency definitions are created. For example, a policy might state that evidence records may be created only if a suitable reusable competency definition already exists in the enterprise's catalog or in a catalog the enterprise subscribes to. Another policy might state that HR employees may create new reusable competency definitions as they analyze various sources of competency data, such as resumes, hiring requests from managers, existing competency or task models, etc.

Ideally, existing reusable competency definitions should be used wherever they are available, rather than inventing new ones. A full exploration of this topic is outside the scope of this paper.

Reporting and analytics

Reports

The proposed model provides a way to generate useful reports that do not require expensive customization, because of the regularity of the data structures regardless of the level or type of competencies being tracked and reported. The reusable competency definition standard and the use of simple hierarchical maps rather than massive or very complex ontologies simplify the

design of comprehensive reports with useful, human readable information. For example, given a reusable competency definition identifier, the title and possibly description of the competency definition can be looked up for inclusion in the report; a hierarchical competency map provides a standard blueprint to structure the report to show component competencies and/or facets of competency and how they "roll up".

Analytics

There is now a way to support meaningful analytics because the data models are coherent and the data to support analytics are unambiguous. Without this model one is forced to try to analyze widely disparate kinds of data with wildly different semantics, or much poorer data that does not adequately capture the variety of processes and sources of competency evidence.

Real-time reporting and analytics are supported if the implementation supports distillation in real time with automatic application of the confidence policies.

Examples of analytics exploiting the proposed model:

- If reusable competency definitions and/or competency maps exist to specify the competencies required to support performance on specific tasks or in particular positions, actual performance measures for an individual or team can be correlated with competency records for that individual or team. Observed performance can be fed back into the system and distilled into standard evidence records that can be used in longitudinal studies.
- The impact of the use of learning resources or performance support resources targeted to specific competencies can be verified because of the common references to reusable competency definitions are the key to relate use and result.
- The effect of different forms of interventions can be compared. Each evidence record contains a pointer back to the raw data that result form the intervention. The effect of the intervention on different individuals or groups can be measured by comparing evidence records that can be traced to the interventions, and by excluding from the analysis the individuals for which other evidence records indicate extraneous influences.
- The confidence ratings can be used in various ways to analyze credibility factors or to filter out unreliable data.

Security, confidentiality and privacy

Authentication and authorization

Authentication and authorization are important to avoid corruption of competency evidence and competency records. The proposed model does not specify how to deal with such security issues. However, it simplifies the security issues by isolating the data into well defined silos with well defined propagation paths.

Intentional or accidental data corruption

The sad reality is that people cheat and that even when they don't we have to work with often incomplete or unreliable data. The proposed model facilitates the implementation of policies to minimize the impact of cheating or bad data. It specifies a simple process to distill the data, minimizing the ways to circumvent sanity checks.

A fully auditable implementation would create a persistent assessment record (which might just be a simple note) or evidence record for every transaction that might influence competency records.

The model also provides a full audit trail that enables manual and automatic corrections. A manual correction should cause the creation of a persistent evidence record that identifies the person who made the correction. A more stringent policy might require that manual corrections be submitted in the form of an assessment form that gets distilled into an evidence record with the proper confidence rating applied by policy.

Automatic corrections can be implemented because of the way the data is cross-referenced in a resilient way. For example, if it turns out that a particular instructor was inflating grades, all the evidence records based on those grades can be identified because they contain a reference to the source of evidence, and their confidence rating lowered or even set to zero. Then, all the competency records that rely on these evidence records can be corrected to fall back on the next best evidence record, if one is available.

Other complex operations can also be automated to implement new policies or competency reference models. For example if two enterprises merge, and some competency definitions of one enterprise are found to be equivalent to competency definitions of the other enterprise, all the competency records can be updated after creating evidence records that specify the determination of equivalence as the source. This provides an audit trail to explain how competency records may exist for the "new" competency definitions even though there is no direct evidence of assessment for some of the employees. Obviously, the same method could also be used if a company that has a legacy competency model acquires a competency model from a vendor and uses it to remap their existing competency records into the new competency model.

Confidentiality and topic security issues

The actual competency definitions may be a trade secret, a national security secret, or may have to be kept confidential for some reason. The same applies to the assessments and assessment results. The proposed model allows management of competency records and exchange of competency information by means of opaque identifiers, without revealing the actual definitions or assessment instruments or assessment results. For example, in this model an evidence record contains no data that describes a topic or assessment data. It only contains generic values and identifiers that are useless unless you are authorized to look up what the identifiers represent.

There is of course a limit to this, which is that if an identifier is compromised, e.g. if someone publishes what an identifier represents, then the identifier can be used to look up the public information. An analogy is the abuse of social security numbers in the USA. With only your social security number, which is supposed to be only an opaque identifier for private data, banks credit bureaus can collect and sell to this data because it is widely shared among trading partners. This type of policy failure is however outside the scope of a technical standard.

Privacy and HIPAA conformance

The proposed model allows manual and automated management and processing of competency and assessment data without exposing individually identifiable information, as required by HIPAA. Everywhere in the model, individuals are represented by opaque identifiers. Where necessary to secure privacy, the identifiers can be obfuscated by common agreement among trading partners. For example, the "person or group" identifier in a competency record might be an opaque handle that can only be resolved to a true identifier through a resolution system that requires proper credentials.

