CEREBYTE

1 ▶ 2 ▶ 3

This is part 2 in a series of articles on the harvest and reuse of expert knowledge. This article focuses on facilitating the practical application of expert knowledge by non-experts to resolve real challenges. Other articles focus on the harvesting process itself and using DCT to significantly expand the value of performance improvements.

8 Minutes
To Performance Improvement

William Seidman, Ph.D.
Michael McCauley

Executive Summary

A Couple of Quick Questions…

Once you have the experts' secret sauce, what are you going to do with it?

How about using it to get an immediate performance improvement?

Effectively using the secret sauce can produce a significant behavior change and related performance improvement in as little as 8 minutes.

Main Topics

- Non-experts can be enticed to review expert content in just a few minutes through a series of questions that tap into personal motivation
- There is only a 90 second window to convince non-experts to take ownership for applying the expert content and changing their behavior.
    - The non-expert must have *instant* engagement which is created by *instant credibility* (the non-expert looks at the content and immediately recognizes the realism and power of the content) and *instant application* (the non-expert can immediately apply the content to their own situation, creating a significant personal benefit0
    - Digital Coach Technology (DCT) creates instant engagement for the non-expert. It supplies the secret sauce to create instant credibility and the on-screen prompts engage the non-experts' "passion" for their work, stimulating both a cognitive and a emotional response that promotes instant application
- 8 minutes after instant engagement, the non-expert changes their behavior. They lean forward more, focus their eyes differently and speak differently, using the experts' language and concepts. They enter a state of accelerated learning from which they behave in accord with the experts' secret sauce.
- Instant engagement creates an opportunity for sustained performance improvement producing these results:
    - 80% reduction in planning time
    - 50% reduction in training time
    - 30-50% reduction in task performance time

Cerebyte, Inc.
Lake Oswego, OR
t: 888.355.0213
f: 503.624.1642
www.cerebyte.com

8 Minutes
All It Takes to Improve Performance!

William Seidman, Ph.D.
Michael McCauley

First, A Story…

This really happened! At a recent knowledge gathering session, someone asked a "stupid" question:

> *"So, now that we've gathered all this wonderful expert knowledge, what are we going to do with it?"*

A stunned silence ensued.

Rewind to the beginning.

The Mergers and Acquisitions (M&A) expert had consistently and systematically outperformed all the other M&A professionals in the company. Learning his "secret sauce" and aligning the processes of the 200 other M&A specialists with his "secret sauce" could produce millions of dollars in additional profits.

The company had previously tried interviewing him, following him around, and videotaping him in an effort to identify and gather his tacit knowledge, but they were never satisfied that they truly understood the subtle factors that made him so successful. Now, with a new approach to harvesting based on "digital coaching technology (DCT)," they knew they had it.

His secret sauce was tremendous. It included such subtle knowledge as:

- The identification and interpretation of a particular type of phrase uttered by the CEO or CFO of the acquiring company that communicated the "real" reasons for the acquisition
- How he aligned the decision-making process with these real reasons
- His ability to focus in on a single clause in a specific type of contract that typically determined the likelihood of a successful merger or acquisition

They were very excited. And then that pesky question was asked.

> *"So, now that we've gathered all this wonderful expert knowledge, what are we going to do with it?"*

They had been so focused on gathering the secret sauce that they had not thought about what they were going to do with it once they had it. Furthermore, they

intuitively knew that the conventional wisdoms of creating a binder or a training course from the material could not possibly communicate the intensity, sophistication and depth of the knowledge. They wondered: **"Why bother to collect knowledge if there is no plan or means to make it useful to the organization?"**

Unfortunately, the story recalled above is both an actual situation and a standard set of responses that are typical of how many organizations have tried to use expert knowledge to improve performance. Too much attention is spent on trying to *gather expertise and create best practices* while too little attention is focused on enabling less experienced or successful people to *reuse that expert knowledge* to improve their own performance.

This paper presents an approach for facilitating the application of expert knowledge by non-experts to real situations. Specifically, it discusses the factors (including DCT) that promote the use of expert knowledge by creating the experience of personal mentoring for many non-experts simultaneously and at low cost. Where this approach has been used, it has caused significant behavioral change and improvement of non-experts in as little as eight minutes. For one fast food chain, the improvements are expected to increase sales by $2,000 per week per restaurant!

A Surprising Omission

Why do organizations place so much emphasis on collecting expert knowledge, and almost none on using expert knowledge? Simply put, collection is the easy part – it's the reuse that is hard!

Getting non-experts to use expert knowledge requires a solid understanding of four separate disciplines, as well as the ability to integrate those disciplines (Figure 1).

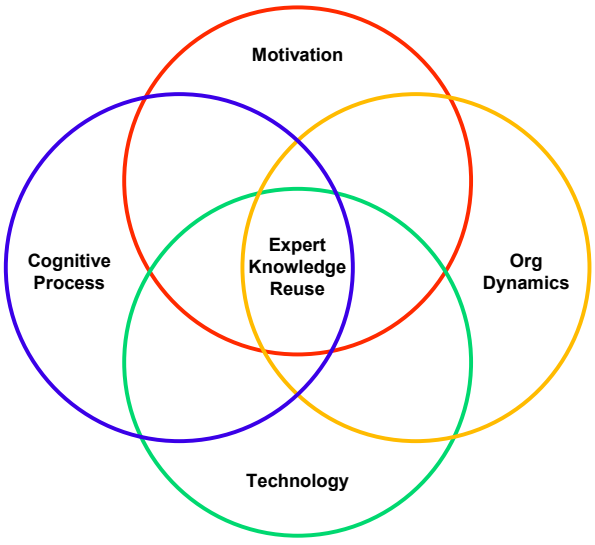

Figure 1: The Zone of Knowledge-based Performance Improvement

The disciplines are:

- Human cognitive processing (how we think and process information) – this is essential for understanding both how the experts structure their knowledge and how non-experts absorb that knowledge

- Human motivation (what makes use act in certain ways)– getting the non-expert motivated enough to take ownership for absorbing the expert content is the single most critical, and most difficult, component of success

- Organizational dynamics (how organizations behave) – both cognition and motivation must be wrapped in an organizational context. If these do not align, no behavioral change takes place

- Technology – this includes both the design of the interface, which is the single most critical technical area, and the underlying architecture of the system that enables the integration of the above processes

As a non-expert attempts to use expert content, these factors interact with each other. For example, guiding the non-expert to take ownership (motivation) for the expert content (cognitive), the architecture of any support system must have a capability for local adaptation and connectivity (technical), as well as a means of centrally storing and maintaining the approved best practices (organizational). All four aspects of using expert knowledge are constantly interacting and must be treated as completely integrated to create quick, effective behavior change.

Not surprisingly, this is not a comfort zone for many people or organizations. Historically, training courses and binders have focused more on the cognitive while e-learning and knowledge management have focused more on the technology. Few, if any, systems are focused on either motivation or organizational factors. To make matters worse, until now, nobody has recognized the importance of an integrated approach.

A Critical Definition

Non-expert use of expert knowledge has received so little attention that there isn't even a commonly accepted definition of the process. Here is our definition:

> *Reuse of knowledge is the application of another person's knowledge to a new situation, without them present, in a way that produces a significant productivity gain.*

There are several key elements to this definition:

- Non-experts must actually use the knowledge in a new situation. This is not a frivolous idea. It is impossible to establish any meaningful value if the non-experts don't use the expert content in a new situation. Yet, the majority of knowledge management training and process engineering is about creating databases, binders and

passive learning systems that produce little or no actual reuse of the knowledge

- They must use the knowledge without the expert present. Again, this is not a frivolous idea. Communities of practice and personnel profilers all require the expert to be present, undermining expert productivity and disenfranchising the experts over time. Reusing knowledge is only meaningful when it is successfully done without the experts present.

- Use of the knowledge must create a "significant productivity gain," otherwise it is a waste of resources. This is the payoff. Knowledge doesn't have value to the organization in itself. It only has value if it produces value. A remarkable amount of "information" has remarkably little sustained value. If you can't foresee a well-defined gain, you should not be collecting the content!

This definition creates a very high standard of performance, but anything less is meaningless. Even though each element of it is essential to knowledge reuse, it will no doubt frighten many people.

Models of Knowledge Reuse

There have, of course, been many previous efforts to use expert knowledge. Generally, these efforts fall into just three categories (Figure 2), each of which has some significant weakness. The three categories are: print media, classroom training and personal mentoring.

Print media, such as a book and the ever-present corporate binder, has the advantage of being relatively inexpensive to expand for additional usage. Once you have written a book or created a binder, it is relatively easy and inexpensive to duplicate it for another person. But these media are expensive to initially develop, tend to be difficult to adapt to unique circumstances, and perhaps most importantly, the application of the knowledge is solely the responsibility of the end user. Actual application and utilization of the knowledge is problematic. Also, print media tends to reflect the company perspective, or the "official story" (e.g., they way things should be done *in theory*), leaving little room for people's reality. Costs are low, but so is the impact.

Training covers a wide range of activities from lectures to workshops. Lectures have many of the same characteristics as print media. Adding another person to a lecture is not expensive, but the initial cost of the lecturer can be substantial. Like print media, the attendee must apply the knowledge to their situation without direct support. Again, overall costs as well as impact are low.

In workshops, these ratios change somewhat. While there is a greater impact of the knowledge due to the direct work of applying it to specific situations, it is difficult and expensive to add additional people to the workshop. For workshops, both costs and impacts are moderate.

Category	Pro	Con
Print Materials	• Low cost of distribution • Can be centrally managed and maintained	• High cost of development • Difficult to adapt • Focuses on "official story" rather than reality • Application of knowledge is responsibility of end user • Limited impact
Training	• Moderate cost of distribution • Classroom application of knowledge facilitated by session leader	• Moderate cost of development • Ultimate application of knowledge is responsibility of end user • Moderate impact
Mentoring	• High, immediate impact • Application of knowledge "coached" by mentor • Low cost of development	• High cost of distribution • Limitations on speed of distribution

Figure 2: Categories of Knowledge Reuse

Mentoring is at the other end of the continuum from print materials. The impact of mentoring is quite high. In fact, when people are asked to identify their main mechanism for reusing knowledge, after tepidly mentioning the above approaches, they invariably state that they call a mentor when something really matters. Mentoring has immediate impact and applicability, but it is very expensive. A mentor can only transfer knowledge to a few people at a time. While the impact is high, so are the costs.

Clearly, the goal in transferring knowledge should be to achieve high impact at low cost. Using the above categories as a model, it would be like achieving the impact of mentoring at the cost of printing. While this is a laudable goal, it is not easy to do because successful mentors do many subtle things to enhance their mentoring impact. In particular, they must immediately motivate naïve people to initially participate and then create almost instant, but sustainable engagement by the non-expert.

It is really the immediate, sustainable engagement that produces the initial performance improvement and, as the above discussion demonstrates, there hasn't been a fully satisfying model for reusing someone else's knowledge…at least until now!

Effectively applied, DCT can produce that initial performance improvement in just 8 minutes!

Motivating Participation by Naïve Personnel

One challenge that previous solutions have had is getting the non-experts to look at the content in the first place. Such resistance takes many forms including: "We're too busy" or "Our culture would never be this organized."

However, on closer examination, it is clear that systematically using expert knowledge is not, in fact, a wholly new activity. People spend considerable time and effort getting organized for any given function they are expected to perform, particularly if they are new to the function. As part of this, they spend time and effort trying to learn from others. Thus, a systematic approach to using expert knowledge to get better organized is just taking a process that is occurring frequently but inefficiently and making it more efficient. Recognition of the opportunity for increased efficiency motivates initial participation by non-experts.

More specifically, initial participation is accomplished through a simple process that leads to a strong desire to use the expert knowledge. Here are some of the questions that motivate participation:

- Do you expect to spend at least a little time getting organized before you perform this function? (Almost always, the answer is "Yes")
- What is the minimum time you expect to spend getting organized for this function? (Typically "several days," if the function is substantive)
- If you could become significantly better organized in just 30 minutes, would you be willing to use a different approach to achieve that result? (Always, "You bet!")

These three questions lead to a simple, final question that almost always results in participation:

- Do you want to continue to do what you are already doing in an ad hoc way or do you want to do it more efficiently? (Again, the answer is always, "If I had the option, I'd rather do it quickly and move on!")

Most people (but, surprisingly not all) opt for improved efficiency and become engaged!

Instant Engagement

Ninety seconds!!! Our experience indicates that, once you get a non-expert to participate, you only have an additional ninety seconds to convince them to continue their engagement. Non-experts are so inundated with information that they have become extremely skilled at screening out anything they don't think is valuable and relevant – immediately!

To effectively use the ninety-second window an approach must have two elements:

- *Instant credibility* for the content – the non-expert must instantly perceive that the content is rich, realistic and useful

- *Instant application* of the content to their actual situation – the non-expert must immediately move from thinking that this is good information to thinking that using the information can personally help them

Instant Credibility

Credibility is largely created by the energy, passion and realism of the content. Does this sound familiar? The foundation of instant credibility is, of course, the type of knowledge harvested using the approach described in the first article in this series. Content that demonstrates the emotional reasons people commit to excellence and its realistic application to varying complex circumstances entices and engages the non-expert.

For example, the more effective managers of a pharmacy chain thought of their job as "Making the pharmacy a fundamental part of the each family's medical support system." Less-effective managers thought of their job as "Writing 120 prescriptions a day." Similarly, top performing managers at a semi-conductor manufacturing company thought of their performance management system as "refreshing and energizing the organization by identifying and rewarding the future leaders" while less effective managers thought of it as being primarily for "equitable pay administration." Which would you rather do, help families and "energize the organization," or write prescriptions and administer pay? Not surprisingly, most new managers are more engaged by the former which leads to a significant, immediate increase in interest and participation.

Conversely, content based on "official stories," as found in most training materials and process binders, is immediately rejected. You *must* have great, realistic content to get any significant performance improvement!

Instant Application

Instant credibility alone does not improve performance, however. Within that first 90 seconds, the non-expert must also be convinced that the content is useful to them personally. They must think of it as somehow improving their lives, making their job easier, or improving their efficiency. This is accomplished by having the non-expert concisely modify the expert content to "fit their particular situation" and reflect *"their passion"* for the subject.

This approach has the effect of first engaging the non-expert in the material, then causing a subtle transition from a cognitive process to a motivational process. Having the non-expert modify the content causes them to carefully read it. The emphasis on being concise causes the non-expert to ask himself, "What is important?" The prompt for passion causes the non-expert to expand that

question and ask himself, “What is important *to me*?” As soon as the “to me” part is added, you have substantial, sustainable engagement.

For example, in the semi-conductor performance management system, non-experts tended to begin adaptation by highlighting and preparing to delete the “refresh and energize” portion of the description. After all, the program was about pay administration! But they always stopped before completing the deletion, gave the system a contemplative look and made comments like: “This is really different, but better than I was thinking” or “This is a really good way to think about the program.” When there was a manager present, an intense dialogue ensued about the focus of the program that produced in-depth alignment by the non-expert with the program goals as well as a motivated non-expert.

Typically, this approach results in almost immediate, observable changes in non-expert behavior. The non-expert leans forward more, focuses more intently on the screen and breathes more rapidly. When there is someone else present, the language of the dialogue becomes more intense and specific to applying the content to the non-expert’s situation. Non-expert behavior is changed, and performance improvement begins, usually in just eight minutes.

Digital Coach Technology (DCT)

However, while these results are impressive, for most people, mentoring is an intensely personal, expensive, one-one experience. How then can an organization possibly achieve these mentor-like performance improvements for hundreds of people simultaneously? Enter digital coaching technology!

DCT is software that emulates the human mentoring experience. It uses a highly structured process (see Figure 3) consisting of a specialized interview that causes the expert to tell his story (called the “naïve new person story”), polish that story into a repeatable best practice, and store the best practice in an archive. When an inexperienced person needs to perform the same function, instead of asking the expert to mentor them, they recall the best practice from the archive and are “coached” through performing the function as though the human expert was there.

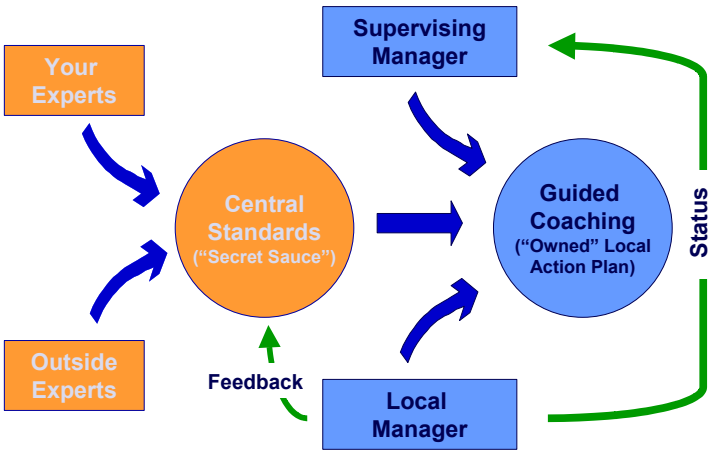

Figure 3: Digital Coaching Technology (DCT) Structure

The actual digital "coaching" is accomplished by a series of on-screen prompts that simulate the questions, feedback and direction provided by an experienced mentor guiding a non-expert to improved performance. The prompts include explicit directions such as "Modify the expert content to fit your specific situation" and subtle screen and system design that guides behavioral change. For example, some DCT fields are deliberately undersized. The resulting subtle visual pressure to be concise causes the non-expert to focus in on what is important to them (one of the criteria for instant application), causing them to take ownership of the content. Similarly, in a section on risks, if the non-expert does not respond to the expert information about risk management, DCT generates a "guilt" screen that causes the non-expert to confront an area of discomfort. DCT combines cognitive, motivational, organizational and technological elements in a way that produces results very similar to the human mentoring experience.

In fact, non-experts typically stop distinguishing between the digital coach and a human coach by the second screen. Thus, DCT creates the impact of mentoring at the cost of replicating software.

Furthermore, because DCT is based on almost universal human capabilities, it can be applied to any specific content. It has been used for such diverse purposes as coaching performance improvement for managing fast food restaurants (think of this as high school kids in hair nets), safety management at naval shipyards and launching interplanetary vehicles.

Sustaining Performance Improvement

Eight minutes to visible performance improvement tends to motivate non-experts. They feel good about the capabilities they can see developing. But immediate improvement does not, necessarily, lead to sustained improvement.

Mentors use several other processes to sustain their initial impact, including:

- Separating the content into a high-level overview and detailed content in a way that enables the non-expert to first develop a sustainable, high-level conceptual framework (a "platform of trust") and then drill down for the application of the detailed information

- Utilizing a sequence of "Read, Feedback, and Adapt" to engage the non-expert in specific discussions about the content and how it applies to their situation

- Driving for commitments to actions, complete with dates, which is done to ensure that the non-expert both truly understands the content and will actively use it

- Continuously monitoring the performance of the non-expert, providing support and further guidance as needed

Together, these processes, again supported by DCT, create both sustained performance improvement and continuous development of the expert content.

The process for creating these is beyond the scope of this article, and will be the focus of the third article in this series.

Alternatives to DCT

Are there viable alternatives to the use of DCT for the application of expert knowledge? We don't think so – at least not today!

Many things have been tried in the recent past, but few have succeeded. The most common past approaches have been process binders, training courses, e-learning, centralized knowledge, and document databases. These have failed for many reasons, including:

- They provide little that motivates the initial non-expert participation
- Their content is consistently rejected by users because it is too "official"
- Their content is often lacks context, becoming data and document trivia
- Their content is usually too generic to be meaningful
- They provide little (or no) emphasis on applying the knowledge (much less a complete tracking capability)
- They are considerably more labor intensive and expensive to create and use than DCT

More generally, these approaches have little understanding of the importance of integrating cognitive, motivational, organizational and technological processes. At best, they may account for two of these elements, but never all four. Yet all four are essential to create effective application of expert knowledge by non-experts.

Benefits of the Applying Expert knowledge

Does this matter? Only if you are interested in these types of results:

- 80% reduction in the time it takes a non-expert to get organized to perform an expert function
- 50% reduction in the time it takes to train the non-expert to perform the expert function
- 30-50% reduction in the time it takes the non-expert to perform the expert function

Translating this into something more specific, here are some of the impacts reported by businesses that have used DCT:

- A global test and measurement company reduced the time required to get organized for the launch of a new product from six weeks to just four days

- A semi-conductor manufacturer eliminated eighteen days per manager per year of repetitive activity

- A shop floor automation software company was able to support nine implementations simultaneously, with the staffing level that previously only supported three

In addition, organizations that already sell knowledge in the form of consulting services have created supplemental product lines by selling knowledge in DCT formats. One safety consulting company expects to double its revenues (from $2M to $4M) by adding a digital safety coach to their offerings.

Summary

The impact of gathering expert knowledge is zero, but the impact of reusing expert knowledge is staggering. Upon reflection, such impacts are not really surprising. Much of what we actually do is based on a foundation of repeatable behaviors. By reusing knowledge, we actually get to rationalize the repeated portions of our behavior to the point where they are very efficient, thereby creating a net productivity gain. DCT can produce the impact of mentoring at the cost of print. It reduces systems costs while increasing productivity – and isn't that what we really want?