


Just in Time KM

Conceptual roots and pragmatic interventions

David Snowden
The Cynefin Centre






CC creative commons
COMMONS DEED

Attribution-NonCommercial-NoDerivs 2.0

You are free:

- to copy, distribute, display, and perform the work


Under the following conditions:

-  **BY:** **Attribution.** You must give the original author credit.
-  **Noncommercial.** You may not use this work for commercial purposes.
-  **No Derivative Works.** You may not alter, transform, or build upon this work.

- For any reuse or distribution, you must make clear to others the license terms of this work.
- Any of these conditions can be waived if you get permission from the copyright holder.

Your fair use and other rights are in no way affected by the above.

This is a human-readable summary of the [Legal Code \(the full license\)](#).

[Disclaimer](#) 

The full copy of the CC licence is available www.cynefin.net – knowledge base.

Just in Time KM

Conceptual roots and pragmatic interventions

David Snowden
Founder
The Cynefin Centre
www.cynefin.net

I first used the phrase Just-in-Time Knowledge Management (JIT-KM) on a conference platform back in 1999 in the context of some recently completed research into the balance of knowledge held in informal and formal communities. I doubt if I was the first and know that I will not be the last. Tom Davenport and John Glaser in an excellent HBR article¹ talk about the application of JIT principles in the context of “embed it (knowledge) into the technology that knowledge workers use to do their jobs”; the phrase is in increasingly common use. One reason for this is the increasing recognition that approaches to knowledge management based on the codification of knowledge to databases that operate on a pull basis have largely failed. Davenport and Glaser recognizes this stating that the focus on networks and communities of practice required knowledge workers to participation in tasks “in addition to doing their regular job. That meant staying a little later each day to share what they’d learned in the course of doing their jobs and coming in a little earlier each morning to learn from others. As a result the programs, many of which continue today, have been only marginally successful”. I came to a similar conclusion, but from a different perspective in the article that reported the above referenced work on informal and formal communities².

“It is not possible to build an intellectual capital management system (ICMS) if we see it as a universal application. Intellectual capital is too diverse, too complex and too heavily dependent on individuals and communities who do not behave rationally. Neither do we want them to behave rationally – to do so would drive out innovation and relationships, both of which are skills for the knowledge economy.”

For Davenport and Glaser the future rests in “bak(ing) specialized knowledge into the jobs of skilled workers – to make the knowledge so readily accessible that it can’t be avoided.” This is not a revolutionary idea; it has been at the heart of operational knowledge management since the early days of knowledge engineering many decades ago and is a context specific activity, which reduces the claim that “this method could revolutionize knowledge management”. Davenport and Glaser base their thesis on clinical knowledge, which deals with a finite set of rules and properties and where the overriding critically of patient care provides focus. This is often not the case in organisational knowledge management projects; Weick and Sutcliffe³ are in danger of making a similar error in assuming that practices common in high resilience organisations such as fire fighting crews and aircraft carriers can transfer into organisation environments where there is not a similar single uniting principle or function. In knowledge management context is the be all and end all of practice. As Davenport and Glaser acknowledge “embedding knowledge into everyday work processes is time consuming and expensive. It’s not an undertaking that anyone in his right mind would tackle without good reason.”

Embedding knowledge is one aspect of sound practice⁴. In this article I want to agree with Davenport and Glaser within the context of their research, but will argue that we cannot universalise from that model into other contexts and secondly that there is a huge area of JIT-KM that can reap a significant return for minimal investment if we focus on accelerating natural processes rather than attempting to impose a “rational” model onto human systems. This complements the expensive but necessary task of embedding knowledge and predates it as an idea; it should also come first to prevent wasted money on embedding knowledge where wither that knowledge is not susceptible to process, or

is heavily time-context dependent. A major problem is that the dominant ideology (and it is an ideology) of management science assumes a constrained set of so called rational practice based on a mechanical design approach that effectively precludes the stimulation of natural forces. This article will therefore deal with the negative before proceeding to the positive.

The false promise of mechanical design

There are at least three fundamental errors at the heart of most theory and practice in Knowledge Management, recognition of which is critical to making progress. My intention is to look at each of these in term, taking a slightly extreme position in each case to make a point against prevailing orthodoxy. The consequence is that we have to take a new look at management of knowledge, which recognises that control through pre-determined goals, and prescribed behaviour has a very limited application. The final section of this paper provides summary examples of JIT-KM interventions that do work, but on the basis of triggering ecological change in which direction can be managed but not the destination. In the conclusion I will correct the balance to a degree arguing that too many people now approaching knowledge from an ecological perspective are naïve in that they fail to recognise the capacity of human society to structure **some** aspects of its inaction in such a way as to permit goal based deterministic management.

THE FIRST ERROR: *It is assumed that, faced with a choice between one or more alternatives the individual human actor will make a "rational" decision based on either minimising pain, or maximising reward. This is not to deny that such targets influence behaviour or that they should necessarily be abandoned, but it is to deny that there is a fully manageable cause and effect relationship.*

Humans, individually and collectively work on the basis of contextual pattern recognition, often at a non-conscious level⁵. Visually only .01% of our visual range is in sharp focus at any time, and we see through multiple point observations filling in the gaps based on previous experience. Skilled craftsmen train hand-brain co-ordination in the same way. The same happens conceptually. Those patterns are entrained based on our own past experience, the collective experience of our culture, often communicated through stories, national and organisational. If the hero stories of our early days in work are of people who held to ethical principles despite the loss of personal reward, then the patterns influence our subsequent behaviour; in contrast a culture based on maximising personal return has inevitable consequences for ethics. Now humans can overcome this pattern entrainment, it is a distinguishing feature from animals but it is not commonplace. Some key patterns, such as those based on trust are particularly difficult to influence or direct, as trust is built over years and lost in seconds, it cannot be trained or enforced.

This patterning of our existence has an obvious impact on decision making with implications for reward and recognition systems; in practice most incentives are threatened punishments and play to the negative but also fail to take advantage of existing patterns before they crudely attempt to create new ones. In practice, within existing trusted relationships knowledge sharing takes place naturally and without inhibition; utilising this is at the heart of JIT-KM. Attempts to compel behaviour through threats and promises tend to induce one of two behavioural patterns:

1. Camouflage behaviour in which knowledge appears to be shared, but is shared in such a manner that it can only be used by reference to the knowledge holder who can then perform a trust validation on the potential knowledge user.
2. Conformance, in which time pressure means the minimum possible, is done to achieve a measurable goal or target. This is more dangerous than camouflage in that it can lead to false confidence that the knowledge has been fully captured.

It is impossible to measure whether someone is sharing their knowledge, but it is possible to measure if they comply with a process.

The mechanisms for process management in consequence do not work for knowledge management. There is a growing body of evidence that camouflage and conformance are more common than full knowledge sharing even when their appears to be measurable compliance; humans are to use that wonderful Scottish word *canny* and are not susceptible to directive control in respect of intangible assets such as knowledge.

THE SECOND ERROR: *Organisations believe that when an expert or other knowledge holder explains how they know something they are disclosing all that they know; and further that the way a knowledge holder says that they know things corresponds with the way that they know things in the field.*

This is a variation of the pattern entrainment referenced above. A simple story will illustrate it from early work using some of the principles of anthropology in knowledge disclosure. The project involved tracking a "trouble shooter" in the oil industry over a week, which involved a scary visit to an offshore rig under fairly extreme conditions. In bad weather a rig is a very nosy place; wind, the waves, helicopters and the rig itself all conspire to create a cacophony of sound. The experience of the helicopter landing itself was disconcerting, but turning (and shouting) to the trouble shooter to ask how we would go about solving the problem that had brought us to this wild place was more so. He said "I already know, can't you hear it?" To me it was noise, but to the entrained patterns of an experts mind he could detect in what he called "the music of the rig" the nature of the problem. By asking how he knew, in the context of his knowing over a dozen heuristics or rules of thumb were revealed. In contrast, asking him about problem solving in the reflective environment of an office in Aberdeen produced a logical, rational process of decision making which more no relation to the reality in the field. If I had created a KM system based on the results of that interview I would have wondered for years why the real experts didn't use it, despite my attempts to involve them in its creation.

The nature of human knowing in all but the most simple of situations cannot be captured in simple models: think of the many failures of artificial intelligence to replace clinical diagnosis. The embedding of knowledge in the structured process described by Davenport and Glaser is an invaluable, if expensive augmentation of human intelligence, but as Davenport and Glaser are careful to point out, it cannot replace that intelligence. If I go to a hospital I want to check the record of the surgeon in respect of a particular operation; at a subconscious level I know that I want evidence that the particular surgeon is not operating from the manual, but is informed by heuristics developed over multiple experiences. In consequence deployment of deep expertise is beyond the capabilities of scheduling systems or the recipe books favoured by too many consultants in knowledge management and organisational change initiatives.

THE THIRD ERROR: *The assumption that a successful model of behaviour discovered in one situation can be understood in terms of its cause and effect relationships in order to create a model of best practice that can be translated, and frequently mandated into another situation.*

Again this error is supported by the arguments above in respect of patterning but can also be understood by making a fundamental distinction between two types of system.

1. An **ordered system**, in which cause and effect relationships exist, can be empirically discovered, verified and which, critically repeat in predictable ways. Such relationships are either known, in which case we can mandate process or are knowable within both an acceptable time scale and resource allocation. This is the dominant system assumption of management science, including most early KM.
2. A **complex system** here, while cause and effect relationships exist they are constantly changing and have many relationships both within and without the system. Moreover in human complex systems the nature of identity adds to uncertainty – humans assume and switch unconsciously between many identities without even thinking about it.

Such systems can only be coherent in retrospect, once something has happened we can explain why, but we cannot predict the future. Think of Betamax against Videotext as standards for video players, any recession or organisational change initiative and you will see the operation of retrospective coherence.

There is a third type, namely a chaotic system, one which is important in modern knowledge management and particularly for innovation but it is not central to the argument of this paper⁶.

We manage ordered systems by understanding of the relationship between cause and effect and then mandating goal-based behaviour, complex systems are very different. We manage complexity by drawing boundaries, to exclude undesirable behaviour or to channel the formation of patterns, and through interventions designed to stimulate the evolution of those patterns into desirable forms. Think of managing children (a metaphor that many a knowledge manager could do with using more often). We draw lines in the sand and say "cross that and you die", draw the line too closely and all authority is lost. We intervene by suggestion or artefact; here is a video, how about a game of football. Only rarely do we attempt to impose a structure and it is now sustainable in other than an emergency. The only difference between adults and children is that adults have become adept at **camouflage** and **conformance** so we may not be aware that our authority is being flouted.

JIT-KM

The time has come to look at positive interventions that can avoid the errors identified above. It is important to emphasise that my overall argument is not that things cannot be managed, but they cannot be managed in the sense that a machine can be managed. One origin of the word *manage* in English is a French word which means the ability to ride a horse in dressage, this is a much more appropriate understanding than command, control and engineering meanings of the process generation.

The following list establishes a set of principles that are not intended to be prescriptive or universal, each of which is illustrated by a current knowledge management practice, some established, some experimental.

Ensure a diversity of approach

Technology is a useful tool, but should not become a fetish. A good tool fits the hand and is largely forgotten despite its usefulness; a bad tool in the technology field too often requires not only the hand, but also the brain, to be bio-reengineered to enable its use.

Mechanical systems tend to conformity and creation of "universal" good practice, as there is in theory an ideal design, organic systems in contrast encourage diversity. One of the main implications of this is to avoid purchasing a knowledge management system in the way that enterprise wide resource planning systems are purchased. Each organisation is a unique context; also individuals work in different ways. Some have messy desks while others are neat and tidy; some enjoy virtual chat others despise it; some multi-task, some work on one task at a time. The rich diversity of human behaviour requires diversity, and attempts to impose uniformity will damage the knowledge base and will not succeed. A knowledge management system is a sound, resilient architecture with many different tools for communication and collaboration in operation. Different communities and individuals can then gravitate to those tools that most naturally support their modus operandi.

In one study in IBM⁷ we identified over 60K private collaboration spaces for just over 50 formal knowledge communities, and those numbers relate only to those who used technology for collaboration so the ratio in practice is even more extreme. Not only is it impossible to formally manage the content of over 60K private virtual spaces, but also it is not desirable.

A fully structured system with taxonomy etc would only discover or manage knowledge that we currently know that we need to know, not that we, which we know or could, know when we need to know it in the future. For example IBM did not know that it needed to know about narrative or story telling skills when it originally designed its formal knowledge spaces, but the existence of a private collaboration meant that the knowledge was self-organised at little or now cost and available when it was needed, in the context of its need. The fact that it was easy and non-intrusive to set up a private collaborative space meant that it was done. Also the selective membership built on existing or newly formed social obligations. I participated in a private space where I had an obligation to fellow "believers" in Narrative long before it was respectable because they were fellow believers, whereas participation in a formal community was a chore.

Don't impact on people's time

Time is the main enemy of knowledge management, followed closely by fear of abuse and at a distance by the more commonly quoted desire for power. In the modern organisation e-mail traffic based on one to one communication has risen to levels where it has a growing negative impact on knowledge worker productivity. We can see this in common e-mail practices, mass copies, blind copying, task avoidance through a request for new data just to get an e-mail "out of the system" and many others. Many aspects of e-mail use are showing all the evidence of addiction. Within the Cynefin Centre we have started to devise programmes in which we use some of the techniques developed in rehabilitation clinics to break addictive patterns of one-to-one communication often by cold turkey techniques: closing e-mail down for a period before bringing up collaborative systems, habituating people to many-to-many tools before allowing one-to-one but then banning all copies and all attachments.

A good general rule to apply is that any knowledge task requiring a time commitment by a knowledge worker requires a prior and explicit gift in the here and now of at least twice that time to allow true co-operation. If you want a group of consultants to devote 2/3 hours a day to knowledge sharing then given them a low paid clerk to handle expenses forms and timesheets – then they will help. Appeal to their sense of duty and loyalty to the organisation and you are doomed.

A developing practice in this field is the use of narrative databases. These use material recorded in the field, which is then accessed in the serendipitous way in which people access key knowledge. Faced with a choice between drawing down best practice from a knowledge management "system" and hearing the stories of 8/9 trusted individuals about their experiences the vast majority would opt for the stories. Narrative works in the same way, recording experiences as they occur and accessing them through high abstraction criteria such as archetypes and themes to reflect the natural process of enquiry. As I enquire of the system I can record new stories in turn and hot link them to more explicit knowledge material. A word of warning though, firstly there are a lot of amateurs operating in the field of organisational story telling (many of them from story telling professions such as journalism who think the skills transfer without amendment) and secondly too many organisations cannot resist telling people what the stories mean and which they should read.

Another interesting approach in this area, which also applies to the next section, is the growing use of apprentice systems, borrowing unashamedly from medieval practice, progressing from apprentice, to journeyman to master. Apprentice schemes evolved to transfer tacit knowledge through observation, coaching and practice and they are not only more effective than manuals and computer based training, but they are also one of the only ways of embedding and validating knowledge transfer between humans. Ironically if the full life cycle costs of knowledge acquisition and deployment are examined they are also generally cheaper than approaches based on rigid codification. Self selecting apprentice schemes can also be created using some of the social network stimulation techniques outlined below.

Connect people, build networks

The number of respected pioneers in knowledge management who emphasise connecting people over codification is only matched by the number of managers who read and praise their thinking then proceed to do the exact opposite, or connect people on the basis of design. We resist the concept of arranged marriages, so why should anyone think the same approach can apply universally to community formation.

To take two examples of new practice in this area:

1. Don't build yellow pages; install some form of expertise locator. Yellow pages rely on people codifying their expertise and maintaining data. In the early days they are novel, so people participate but within a year the system falls into abuse as people either do not have the time, or manipulate their entries based on perceived futures within the organisation they work for. Expertise locators on the other hand are non-intrusive in that they either map affinities based on access to knowledge or trawl e-mail to indicate evidence of expertise. The better ones then respect the paradox of privacy: if I allow someone to keep their knowledge private then they will share it, if I tell them to share it they will keep it private. To be asked in the context of need if I know something is more likely to elicit a response than to be asked to codify in the absence of that context.
2. Social network stimulation (SNS)⁸ works on taking the natural process of network building and accelerates it. Normally over several years as I work in different departments, get allocated or projects and meeting people in a social context I build a network, and hear stories that allow me to operate effectively in an organisation. Mentor schemes, good induction and narrative can all help here, however it is now possible to reduce five years of accidental networking to five months, or even five days in a task based environment. In SNS we focus less on managing knowledge, but on managing the channels through which knowledge flows.

A theme of the above examples is to allow communities to form based on natural preferences, although SNS does it within a set of top down determined heuristics. Telling people that they should work together can often clash with the basic chemistry of human interaction and is thus inefficient.

See where people walk before you build the paths

A good designer observes the pattern of human interaction before they design. For example planting a grass and observing where people walk before you invest in building a path. Of course you can also build a hedge, or create a bridge to encourage patterns: just as we do with children through boundaries and interventions. Here we need to understand that in dealing with complex systems all interventions are also a form of diagnosis and vice-versa; partly because any attempt to study a complex system changes the system being studied. This means that in managing a complex system, and by implication all knowledge management is predominantly about complex rather than ordered systems, we intervene with multiple small probes so that the possible patterning of order is revealed.

Some examples will illustrate this approach:

1. The way we represent human requirements is critical. There is an old adage amongst IT professions that users never really understand what they want by way of system functionality until they get it, at which point they want something different. Equally, many business users are pretty much convinced that the IT profession created user requirement specifications to ensure that they would sign up to something they could only partially understand, but for which they can be held accountable later.⁹ This tendency has been all too present in knowledge management system design. Narrative techniques, particularly those, which represent abstract features of organisational culture through archetypes and

themes, provide a new way of creating a dialogue between system designers and users. Archetype families provide a representation of the different attitudes, experiences and belief systems of a user community. Allowing a designer to put any feature into a system, subject to their being able to create a coherent story about how each archetype will respond to that feature is far richer than esoteric interpretations of a formal document. In effect this creates a fictional space of discourse, which allows ideas to be tested and evaluated in advance of implementation especially if coupled with JAD¹⁰ sessions and prototyping.

2. In forming communities, it is always better to use an existing naturally occurring community, formal or informal where working practices and trust issues are already developed than to impose a utopian design; most KM practice has gone down the utopian design route, which is not to say that it has not worked, but when it has, the cost has been prohibitive. The use of Social Network Analysis techniques, operating between existing communities allows the identification of naturally occurring communities that can be formalised, or clustered based on affinities into cohesive communities. Where a naturally occurring community or cluster cannot be identified then an insect model can be used: swarming. Here the metaphorical equivalent of a bright light is created and we see who swarms, a cohesive entity that results then becomes out new community. Clustering and swarming should be able to provide the bulk of organisations requirements for communities of practice. The other interesting result of this work is that different solutions are required for similar groups within the same organisation, because even within the context of that organisation their development has taken place in a different context.
3. Narrative databases provide a quick and easy way in which to see how a lessons learnt programme should be implemented, and here the experiment may end up as the solution. Any knowledge holder can record to a tape recorder in two to three minutes what they would otherwise take two to three months to spend a few hours writing up. The sheer volume of material that can be captured in narrative form is far higher than can ever be formally codified. Because such systems are self-indexing they can be built and populated quickly allowing the central designer to observe patterns of use before committing to the cost of codification. The designer can see which stories are accessed in what context and how they are used in practice, rather than trying to hypothesis on the basis of interviews which do not reflect the real nature of acts of knowing in the first place.

Conclusion

All of the above examples rest on observing the natural process of knowledge creation and transfer and then use technology as a tool, if appropriate to ensure that such natural processes can be accelerated and scaled. In all cases they focus on self-organisation, all be it within imposed boundaries. In complex systems we embark on journeys and respond to the environment, in ordered systems we determine our goal and then engineer to achieve it. The results and the approaches are different and complementary.

We also need to understand that the balance between these will change over time. When I first probe a space I may use SNS and a narrative database to reveal the nature and possibilities for knowledge use. That experience will allow us to determine the appropriate balance between the techniques identified above and more formal approaches to codification. A unique feature of human complex systems is our ability to create stable patterns of order than allow us to behave as if a system was ordered. This is one of the reasons to reject naïve applications of complexity science derived from, for example, ant behaviour. Humans have free will to create new order, conform to existing order and to disrupt established order. What is needed is an appropriate balance

between each of these, and awareness of the boundary conditions that make current or so called "best practice" dangerous.

JIT-KM is not just about embedding knowledge into process through artefacts, although that is valuable and often necessary but expensive. It is also about creating a human ecology in which knowledge self-patterns, and being aware of when those patterns are stable enough to justify the cost of embedding. It is about diversity within boundaries, using technology as a supporting tool, not as an overriding fetish and above all about recognising that humans are more than capable of achieving results without constant control. Very few senior executives would dream of treating themselves or their children in the way that they treat their employees and it is time for less hypocrisy in system design.

Notes

¹ **Davenport, T.H. & Glaser, J** "Just in Time Delivery comes to Knowledge Management" in the *Harvard Business Review* July 2002 pp 107-111

² **Snowden, D. (1999a)** "Liberating Knowledge" Introductory chapter to *Liberating Knowledge* CBI Business Guide, Caspian Publishing October 1999

³ **Weick, K.E. & Sutcliffe, K.M.** *Managing the Unexpected* Jossey-Bass 2001

⁴ I am using "sound practice" to replace "good practice" which is a term I would ban from KM; good practice is nearly always entrained past practice.

⁵ There is a growing range of interesting work taking place in this sector. Gary Klein's book *Sources of power: how people make decisions* should be mandated reading for executives and knowledge managers; also Andy Clark's *Being There: Putting Brain, Body, and the World Together Again*: both are published by MIT Press.

⁶ The nature of these three systems, or ontologies and their impact on knowledge management is more extensively covered in another article: "Complex Acts of Knowing: Paradox and Descriptive Self Awareness" in the *Journal of Knowledge Management* – Vol. 6, No. 2, (May)

⁷ Published in *Liberating Knowledge* referenced above and updated in *Complex Acts of Knowing*

⁸ A technique paper on SNS is available from the author

⁹ This is a direct quote from a soon to be published article relating the use of archetypes in an advanced KM system design linking methods from both IBM and Microsoft

¹⁰ Joint Application Design, a workshop technique originating in IBM Canada some decades ago and now subject to formal standards through the DSDM consortium.