

Errata

Managing Maintenance Knowledge in the Context of Large Engineering Projects: Theory and Case Study

[Journal of Information & Knowledge Management, Vol. 2, No. 2 (2003)]

William P. Hall

Tenix Defence, Williamstown, Victoria 3016, Australia

and

*School of Information Management & Systems,
Monash University, Caulfield East, Victoria 3145, Australia*

Abstract. Tenix Defence, one of Australia’s largest defence contractors, depends on winning bids and managing contracts for long-lifecycle engineering projects. The ability to capture, manage and deliver project knowledge in explicit formats is crucial to its success. Tenix is moving from a paradigm of traditional paper documents to electronically managing and automating structured knowledge artefacts in a knowledge management framework based on Karl Popper’s (1973) three worlds of knowledge. The new technology captures the authors’ implicit knowledge that was inevitably lost when working with paper documents and also moves aspects of personal cognition from the subjective and personal World 2 into the objective, virtual and persistent World 3.

Keywords: Epistemology; Karl Popper; Michael Polanyi; Explicit knowledge; contexts; annotation.

1. Introduction

Tenix Defence’s implementation of processes and technologies for developing, managing and delivering organisational knowledge — primarily relating to the

operational maintenance of warships in service — is traced from 1990 up to today.¹ I review some contextual imperatives that led Tenix to adopt state-of-the-art content management technologies for building, managing and delivering knowledge-based products such as warship maintenance procedures.

As discussed below, the design of Tenix’s system for managing maintenance knowledge was based on an epistemological paradigm² developed by Karl Popper (1972a), which is somewhat incommensurable (Kuhn, 1962, 1970, 1977, 1983)³ with the paradigm derived from Michael Polanyi’s (1958, 1966) work of largely tacit knowledge that seems to implicitly pervade the Organisational knowledge Management (OKM) discipline. The case illustrates several aspects of Sir Karl Popper’s (1972a) “epistemology without a knowing subject” and three worlds of knowledge, and may help to advertise the importance of Popper’s epistemological framework to the OKM discipline.

¹Details of Tenix and the ANZAC ship project can be found on <http://www.tenix.com>. See also Hall (2001a, b).

²I was a practising biologist for several years, including two years of postdoctoral research in the area of the history and philosophy of science (Hall, 1983) before moving into the area of document and content management.

³Some statistics from Google and ISI’s Web of Knowledge illustrate the evident difficulties authors familiar with one paradigm of knowledge have in discussing the other: Searching Google on 14 March 2003, there are 59,900 references to “Karl Popper” and 6600 to “Michael Polanyi”. Of these only 539 pages (less than 1 per cent) reference both. Of the 539 referencing both authors, 161 include the term “politics” and only 30 include the term “knowledge management”. A similar search using ISI’s Web of Knowledge (as at 12 February 2003: Database(s) = SCI-EXPANDED, SSCI, A&HCI; Timespan = 1997–2003) yielded 913 articles citing Karl Popper’s Objective Knowledge and/or Conjectures and Refutations (ed. and date unspecified); 996 hits for articles citing Michael Polanyi’s Personal Knowledge and/or Tacit Dimension; and only 44 papers citing works by both authors. Of the 20 top-ranked [“knowledge management” Popper Polanyi] pages in the Google searches, only four authors Capurro (2002), Firestone (2000, 2001), Gaines (2003) and Moss (2002) provided any significant comparisons of the two epistemologies. Most of the other hits were lists of books, course outlines, or other topics — not discussions of the book contents.

1.1. *Managing objective knowledge in Karl Popper's world 3*

To date, the relatively new OKM discipline has focussed most of its attention on personal or tacit knowledge (e.g., Nonaka & Takeuchi, 1995; Saint-Onge, 1996; Snowden, 2002; Sveiby, 2000) to the detriment of understanding and managing persistent forms of explicit knowledge in organisations that transcend the involvement of individual people in the organisations. The focus is possibly due to the reliance of pioneers in organisation theory (e.g., Nelson & Winter, 1982) and knowledge management (e.g., Nonaka, 1996; Sveiby, 1994, 1997) on Michael Polanyi's (1958, 1966) works on personal and tacit knowledge, and the fact many practitioners have simply followed the pioneers' concepts of knowledge without ever considering for themselves the theories of knowledge underlying the discipline. For example, in at least some sectors of the discipline one hears the comment "if it's explicit it is only information, it can't be knowledge." A somewhat related criticism has been applied to the whole discipline of knowledge management from another direction: Wilson (2002) argued that knowledge management is simply a faddish name for information management and that the fad is a "bandwagon [that] lacks wheels". In both cases, the problem is that few practitioners have actually looked into the discipline's foundations to understand the fundamental nature of what it is we are trying to manage.

The difficulty here is that there are several theories of knowledge. The most popular are based on or derived from Sir Karl Popper's (1934, 1963, 1972a) works. Popper's ideas have provided an epistemological foundation accepted and elaborated by essentially all the physical and life sciences (excepting behavioural psychology), political science, military affairs, economics, systems theorists and much academic philosophy. In contrast to Popper's work, Michael Polanyi's (1958, 1966) concepts of personal and tacit knowledge are highly subjective, very narrowly based, and have been used significantly only within the domains of religious philosophy, behavioural psychology and management theory including OKM — and this despite the fact that Polanyi was a renowned physical chemist before taking up philosophy (Sheppard, 1999).

The knowledge management system is grounded in a paradigm of knowledge based on Karl Popper's (1972a) three worlds of knowledge. To summarise:

- World 1 is the physical universe, as it actually exists independent of any perceptions, that is, actual truth and reality.
- World 2 is the world of our subjective perceptions, experience and cognition, that is, what knowing entities or

subjects think about the world. Note: Polanyi's (1958, 1966) theory of knowledge is based entirely within World 2.

- World 3 is "*the world of logical contents of books, libraries, computer memories, and suchlike*" (Popper, 1972a: 74 — my italics).

World 2 (our perceptions, consciousness and cognition) is an emergent property of our physical existence in World 1. Personal knowledge and memory form World 2.

World 3 knowledge is created (produced) by World 2 activities, but it has an objective (though it can be virtual) existence completely separate from "knowing" individuals who assemble the knowledge. As well as tangible objects such as books and records, Popper (1972) also explicitly included in his concept of knowledge claims about the world held in computer memories and even species' genetic heredity. Objective knowledge persists independent of any individual and as such can be retrieved back into World 2 by other individuals; or, as the Tenix case will show, aspects of cognitive processing may be externalised and automated. For example, the objectivity of World 3 knowledge is demonstrated by what can be done with the hereditary knowledge encoded in sequences of DNA nucleotides in our chromosomes. Living cells replicate the same knowledge over many generations of cell division such that no one cell has more than a few of the physical atoms of the DNA in the original cell. This is then translated into sequences of RNA nucleotides and then into a sequence of amino acids in proteins. Beyond this, humans can print the coded knowledge on paper and/or send it down a telephone line to a computer for analysis, comparison or even encoding the knowledge back into other DNA elsewhere in the world (U.S. Department of Energy Office of Science, 2003).

As illustrated diagrammatically in Fig. 1, Popper (1972a) assumes that the logical content of World 3 may connect to World 1 by attempting to describe and predict World 1 as it exists, where these claims to know reality can be tested by actions and observations based in World 2.

Writing in terms of scientific theory, Popper posits some important theses about this World 3:

- We can discover new problems in World 3 which were there before they were discovered and before they ever became conscious; that is, before anything corresponding to them appeared in World 2.
- Thus there is a sense in which World 3 is *autonomous*: in this world we can make theoretical discoveries in a similar way to that in which we can make geographical discoveries in World 1.




Fig. 1. Karl Popper's three worlds of knowledge.

- Main thesis: almost all our subjective knowledge (World 2 knowledge) depends on World 3, that is, to say on (at least virtually) *linguistically formulated* theories. *Example:* our “immediate self-consciousness”, or our “knowledge of self”, which is very important, depends very largely on World 3 theories; on our theories about our body and its continued existence when we fall asleep or become unconscious; on our theories of time (its linearity); on our theory that we can pick up our memory of past experiences in various degrees of clarity; and so on. With these theories are connected our expectations of waking up after falling asleep. I propose the thesis that *full consciousness of self* depends upon all these (World 3) theories, and that animals, although capable of feelings, sensations, memory, and thus of consciousness, do not possess the full consciousness of self which is one of the results of human language and the development of the specifically human World 3.

... The commonsense theory of knowledge [i.e., the subjective knowledge of the knowing subject, which Popper (1972a) has discussed at some length] is unaware of World 3, and it thus ignores the existence of knowledge in the objective sense. (Popper, 1972:74 — his emphasis.)

As defined by Popper, this third world is a unique kind of domain:

... I suggest that it is possible to accept the reality or (as it may be called) the autonomy of the third world, and at the same time to admit that the third world originates as a product of human

activity. One can even admit that the third world is man-made and, in a very clear sense, superhuman. *It transcends its makers.* [my emphasis.] [from the footnote] Although man-made, the third world (as I understand this term) is superhuman in that its contents are virtual rather than actual objects of thought, and in the sense that only a finite number of the infinity of virtual objects can ever become actual objects of thought...

That the third world is not a fiction but exists in “reality” will become clear when we consider its tremendous effect on the first world, mediated through the second world. One needs only to think of the impact of electrical power transmission or the atomic theory on our inorganic and organic environment... [Popper, 1972, p. 159.]

The main concern of Popper's work (1934, 1963, 1972a) was to understand how logical processes carried out by “knowing” individuals caused the quantity and epistemic value of World 3 knowledge to grow. Claims to knowledge that are limited in scope and offer few, if any, testable connections to World 1 reality have little value compared to claims that logically relate many assertions about reality that offer many testable connections. The latter kinds of knowledge can readily be improved or eliminated (if demonstrably not useful) by constant criticism of their claims against the reality of World 1. It is obviously beyond the scope of this paper to paraphrase the several books Popper wrote in developing his theory of

knowledge,⁴ but I hope to demonstrate in a small way how the systems that Tenix has implemented help to assimilate and improve the quality of maintenance knowledge held in World 3 formats available to the maintenance organisation comprising Tenix, its alliance partners and fleet operators.

In the discussion that follows, I will also use definitions for three flavours of knowledge derived from those of Nickols (2000):

- *Explicit knowledge* is assessed, aggregated, assimilated and written down (i.e., knowledge that objectively exists in World 3). In Tenix, technical writers assemble and transform input from a variety of sources and their own personal knowledge into World 3 objects. Explicit technical knowledge can be managed by managing documents.
- *Tacit knowledge* refers to natural talent and skills that can be learnt by experience and practice, but cannot be articulated in words, although it may be passed on via mentoring and apprenticeships. Tacit knowledge is personal and is managed by identifying and managing the people who hold it. This is World 2 knowledge that cannot readily be expressed objectively, although with some degree of effort it may be possible to transform tacit into World 3 objects, such as training films and Class-5 interactive electronic technical manuals containing artificial intelligence (DAU, undated).
- *Implicit knowledge* is where objective and personal paradigms of knowledge overlap. Implicit knowledge is held in individual memories and is capable of being made explicit, but has not yet been. Management goals are to identify such sources and develop tools and processes to ensure that it is recorded and made available in the persistent organisational memory. A significant part of the Tenix solution provides the means to readily transfer implicit knowledge into World 3 structures.

1.2. *The life and death and economic importance of World 3 knowledge in World 1 activities*

Even with World 3 documents as seemingly mundane as maintenance procedures, the objectivity of the contained “knowledge” may be tested in life and death situations. If the knowledge does not exist, is incorrect, or is not accessed or used when needed, people can end up dead. For example:

- On 5 May 1988, an engine room fire on board the Navy transport ship *Westralia* resulting from a broken high-pressure fuel line killed four crew members and hospitalised five. Inadequate fuel lines had been installed because maintenance management personnel and suppliers failed to use the assimilated knowledge of established configuration management procedures (Commonwealth of Australia, 1998).
- The 25 September explosion and fire in Esso’s Longford gas plant killed two workers, injured eight, cut gas supplies to the entire state of Victoria for two weeks, and caused millions of people to suffer loss of employment, livelihood, heating and hot water. The loss to business was estimated to be \$A 1.3 billion (Clarke, 2000; Dawson & Brooks, 1999). The explosion resulted when a restored hot oil supply fractured a gas heat exchanger that had frozen because of an earlier failure of the hot oil supply (Longford Royal Commission, 1999). As stated by DNV (2000), the Royal Commission found that “knowledge management at the 30-year-old plant was a key deficiency . . . ‘The lack of knowledge on part of both operators and supervisors was directly attributable to a deficiency in their initial or subsequent training. Not only was their training inadequate, *but there were no current operating procedures to guide them in dealing with the problem which they encountered on 25 September 1998*’”. [Italics mine.]

An important consideration in both cases is that operators can only ever be counted on to have the knowledge required for their day-to-day activities and reasonably common occurrences in their heads. And even then, to train the operators to that state requires explicit training plans, technical manuals and a variety of other explicit knowledge-based materials to guide the training. Beyond their common activities, plant and equipment operators require ready access to a much broader knowledge base than can possibly be held in personal memories in order to deal with infrequent, unusual and extreme events — especially those that have the potential to become catastrophes if not dealt with correctly. For example, some critical ship maintenance tasks (major overhaul of a diesel engine) are performed only once in seven years. Appropriate and correct documentation in a rapidly accessible (e.g., computerised) format can be a lifesaver — or at least essential to keeping the product in service for an economically long life. On the other hand, if documentation is available but is incorrect, people may die or the equipment may be destroyed as a result.

⁴I am exploring the epistemology of individual, scientific and OKM in a book-length hypertext under the working title, “Application holy wars or a new reformation: A fugue on the theory of knowledge”.

1.3. Contextual goals for managing project, engineering and operating knowledge

The organisational imperatives of an engineering project management company like Tenix (the “supplier”) are to win contracts and to profitably complete them. Enhancing the quality of business intelligence and reducing project cycle times are of competitive value. However, to fulfil contracts, besides delivering tangible products, the supplier must also deliver knowledge (i.e., documents) about how to operate and maintain the contracted products. Such knowledge products need to be:

- Correct and consistent for the points of use.
- Applicable/Effective — applicable to the configuration of the individual ship/vehicle and effective for the point in time re-engineering changes, etc.
- Available to those who need it, when and where it is needed,
- Useable, in terms of being readily understandable by humans and, where computer systems are used to manage and deliver knowledge, also readily managed and processed by the computer systems.

Other essential client goals for supplied assets that need to be satisfied by the supplier include the client’s need to minimise acquisition and support costs over the assets’ entire lifecycle.

The supplier’s internal organisational and economic goals for producing the knowledge are to produce high-quality products that satisfy the customer and generate more business, and to do this fast and for a low labour and facilities cost.

2. Project Document and Content Management

The case study reviews Tenix Defence’s knowledge management requirements relating to the ANZAC ship maintenance documentation and some of the technologies we have implemented and developed to manage that knowledge better. Some unusual aspects of the ANZAC ship project have made it more knowledge intensive than many otherwise similar defence projects. The study documents early stages in what is still very much an ongoing process of technological and organisational change to improve our production and management of World 3 knowledge.

Large engineering projects such as building the ANZAC ships or major facilities and plants all have similar lifecycles (Fig. 2). The major deliverable product from most project phases is explicit knowledge in the form of engineering and textual documentation. Only the production/construction phase delivers tangible products.

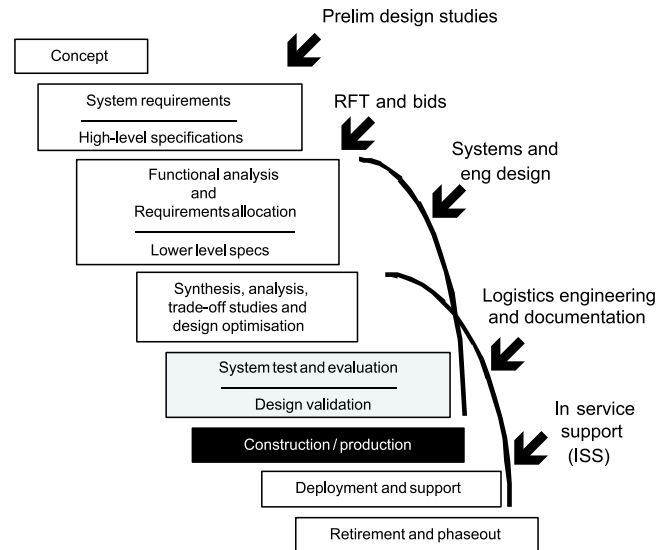


Fig. 2. Blanchard’s stages of the project lifecycle (modified from Blanchard, 1992).

Engineering documentation (e.g., drawings and technical data) is closely tied to the design and production of the tangible products and is well understood by corporate managers (many of whom are engineers). Textual documentation is less well understood by engineering company management, although it is the main bearer of operational knowledge from the supplier of the product to those who must operate and maintain it after delivery.

For large defence projects (e.g., ships, airframes and armoured vehicles), the engineering phases may span a decade, production of the fleet may continue for two or more decades, and the in-service life of delivered assets may span three or four decades. All project phases produce important and valuable textual documentation, without which the project management company or the end-user of the product cannot function. For example, engineering and operating knowledge created early in the project must be preserved, maintained and kept useable for as long as units produced by the project remain in service — which is substantially longer than the survival time of any particular person, system or even many of the technologies used by the support organisation.

Figure 3 shows the flow of textual knowledge through the lifecycle of a major project. In such a project, documents are not static and do not exist in isolation from one another. Specific elements of knowledge (i.e., as reusable text) flow down from one stage to the next with or without modification and may be repeated across many documents at any one stage. In addition to sharing the same elements of knowledge across documents in one project stage or up and down through different stages, the contents of the documents are associated through

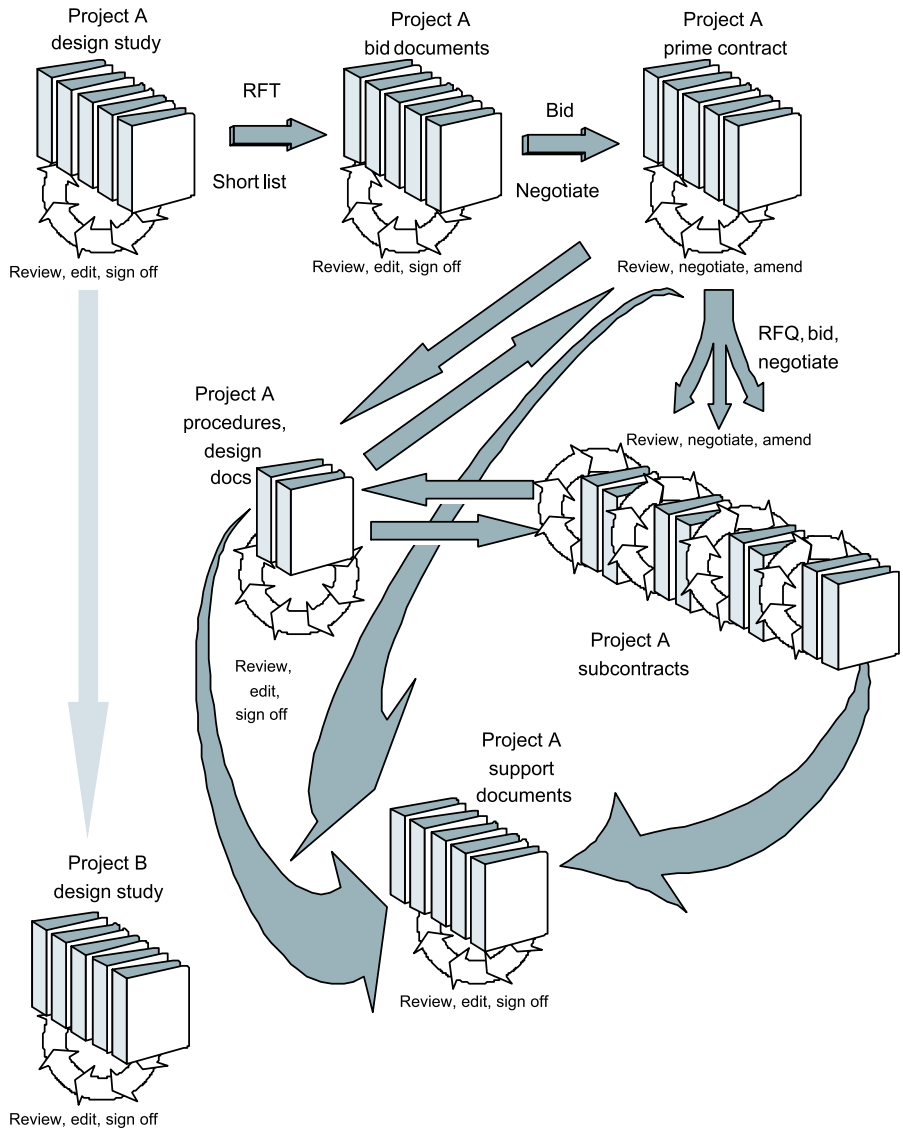


Fig. 3. Flowing explicit knowledge through a large engineering project.

actual and inferred cross references and parent-child derivations into an overall web of World 3 knowledge relating to the product design, engineering, operation and maintenance. When documents are drafted in an electronic environment, elements of knowledge exist and can be managed independent of their representation on paper. Understanding this can lead to major process, productivity and quality gains by reusing or “sharing” this knowledge. In project documents, more than 50 per cent of paragraph-level text exists in more than one place, and in some cases (e.g., in maintenance routines, the same texts need to be replicated across tens of thousands of documents). It may also be necessary to quickly trace the impacts of changes (e.g., to safety information) in one document through the entire knowledge base.

3. Ship Support Knowledge as a Driver for OKM Innovation

3.1. *Special demands of the ANZAC ship project*

The ANZAC project contract as negotiated in the late 1980s is novel in several respects:

- The contract has a stringently fixed price except for currency and inflation adjustments.
- The ships must meet contractually agreed operational availability (“A_o”) thresholds, that is, to be fully available for combat 80 per cent of the time to meet the client’s capability requirements. Fourteen different individual “critical” systems were warranted to be available 90 per cent of the time.

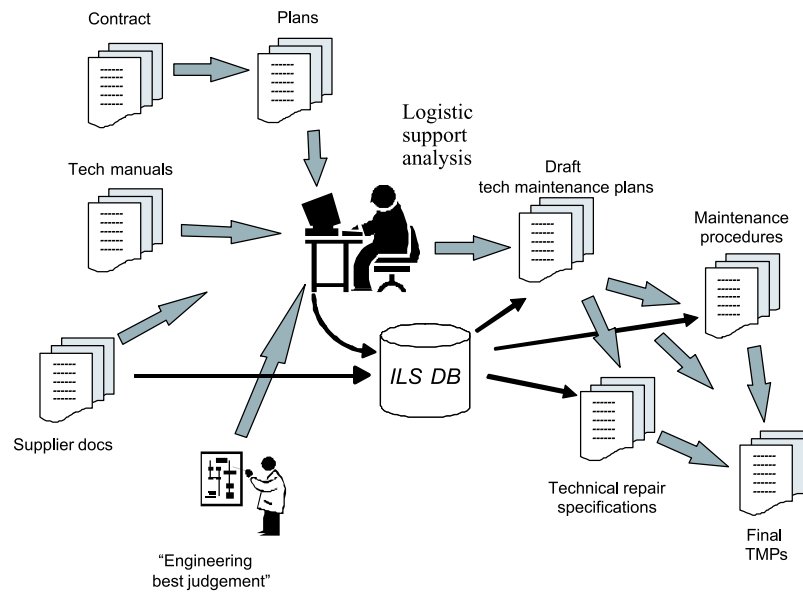


Fig. 4. Assimilation of maintenance knowledge by Tenix into deliverable knowledge products.

- Tenix is contracted to deliver the entire logistic support package, including maintenance documentation, which had previously been the Navy's responsibility, at the fixed price.
- Tenix contracted to develop a Test, Evaluation and Validation (TE&V) process to capture, measure and report the ships' in-service performance and operational availability for the first 10 years of in-service operations⁵ to objectively test whether the knowledge embodied in logistic support package enabled A_o thresholds to be met. Any shortfalls had to be adjusted at contractor expense to meet the agreed A_o thresholds. Tenix developed an Operational Availability Recording and Reporting System (OARRS) to collect and analyse the necessary data. Data collection ended on 19 October 2000, with the ILS TE&V process completed in December 2001.⁶

The 1989 ANZAC ship contract assumed that support and maintenance knowledge would be captured and delivered via paper documents. Tenix's adoption of electronic authoring and the client's adoption of the AMPS computerised maintenance management system⁷ caused a radical change in document paradigms, where the primary delivery changed from physical paper to an elec-

tronic virtual format bearing little relationship to the sequential logic of a paper document.

Scheduled maintenance is critical for the safety of the ships, crews and the continued operational availability of the ships and capabilities. Schedules and instructions encompassed in the documentation deliverables are important knowledge to keep the ships operational in service and protect the health and safety of the crew performing the maintenance, and must reflect specific configuration details of each ship where the maintenance is performed.

3.2. Capturing and distilling maintenance knowledge

Tenix writes and manages all the ship maintenance routines and related logistic support knowledge products, which are significant deliverables under the ANZAC ship contract. Fig. 4 shows conceptually how information is aggregated and assembled by logistics analysts and technical authors into knowledge about how to maintain the components of a ship.

- *Contract.* The contract governs document delivery requirements and standards.

⁵I joined the company in January 1990, one month after the prime contract was signed, and have been closely involved with contracts requirements analysis, system design and implementation of all aspects of the knowledge management system described in the case study.

⁶That is, 4 years for Ship 01, 3 years for Ship 02, 2 years for Ship 03 and 1 year for Ship 04.

⁷Integrated and supplied by Eden Technology, and first known as ANZAC Maintenance + Planning System, AMPS is now being marketed under the name, Asset Managing and Planning System. <http://www.eden.com.au/product/amps.shtm>.

- *High-level plans.* Mandated logistic support plans are developed along with ship and system design activities early in the project. These plans define overall operational, maintenance and crewing policies, which then govern other lower level deliverables in Fig. 4.
- *Draft technical maintenance plans (TMP).* High-level plans guide drafting of preliminary TMPs for each system. A TMP defines a maintenance policy for a system, breaking the system down into individual items of maintenance interest. Finally, for each maintainable component in the system, the TMP sets the schedule and nature of maintenance activities that need to be performed. TMP authors aggregate knowledge from many sources — documents ranging from product brochures and parts lists to formal technical manuals provided by system and equipment suppliers. Authors also use a number of different analytical processes and tools (e.g., level of repair analysis — LORA, life cycle cost — LCC, failure modes effects and criticality analysis — FMECA).
- *Engineering best judgement.* Logistic analysts add personal knowledge gained from years of prior experience in the Defence forces or commercial environments in maintaining and managing similar kinds of equipment, and in the process make explicit some implicit personal knowledge.
- *ILS database (product data management).* Besides assimilating maintenance knowledge into documents, logistics analysts load information about equipment, components and parts relevant to the maintenance activities into an ILS database (ILSDB). Line items in this database are related to the overall systems hierarchy for the product. Configuration sensitive items of knowledge included in the maintenance documents must be maintained for each deliverable against the current configuration of the deliverable product based on engineering master data.
- *Planned maintenance documentation, including maintenance requirement cards (MRCs) and technical repair specifications (TRSs).* There are two classes of planned maintenance activities: those performed by ship or shore personnel on the equipment *in situ*, covered by MRCs; and overhaul tasks done by external repair agents, covered by TRSs. MRCs are either scheduled on a calendar basis or triggered by conditions, such as the number of running hours against a particular piece of equipment, or events, such as preparing for sea. Both classes of documents are drafted by logistics analysts and may be further distilled or rewritten by qualified technical writers. Again, the contained knowledge comprises an amalgamation of maintenance philosophy based on high-level plans, logistics analysis, and information gleaned from documents provided

by equipment suppliers, plus engineering best judgement. In many cases, documents are physically tested against equipment by having someone actually perform the maintenance as described in a draft document to verify that the content is correct and useable. Each ship requires around 2000 MRCs. More than 600 separate TRSs cover external overhaul and repair requirements for the class as a whole.

- *Final TMP.* Many things are learnt in the process of preparing the detailed MRC and TRS documents that reflect back into changes in the system maintenance plans. In other words, both forward and backward knowledge flows are required to build a completely coherent view of the maintenance requirements.
- *Subsidiary extracts and reports.* Although not shown in Fig. 4, many subsidiary reports are extracted from MRC records to provide information required for crewing plans, spares allowances and supply chain holdings, fluids and lubricant requirements, miscellaneous materials, tools, special test equipment, etc.

However, from the point of view of the operating organisation, the most essential knowledge must be distilled into the MRCs and TMPs, because these form the interface between the human maintainers and the bulk of World 3 knowledge about what is required to keep their ships operational.

3.3. *Adapting to changing client delivery requirements*

Because Tenix started drafting maintenance routines in late 1992, before we had a final agreement with the client on delivery formats, we began building MRC content in WordPerfect merge tables to separate the authoring of content from the formatting functions. Because much of the required information could be captured in parsable fields held in semantically defined merge table fields, it was possible to build merge macro processes to extract from the structured content the wide range of reports and internal deliverables we required. By the time the first ship set of MRCs was delivered in 1995–1996, over 20 different outputs were being extracted automatically from the structured knowledge embodied in the MRC records. Also, around this time the client decided to implement AMPS and requested that we deliver MRC content electronically into the relationally based AMPS system. Each MRC required two styles of delivery:

- *“Metadata” to provide scheduling and other job-related information.* This had to be delivered as comma-delimited relationally structured tables to provide scheduling and other job-related information able to be

taken up by the relational database engine at the core of the maintenance management system. This metadata includes configuration identifiers, codes defining how maintenance is scheduled, document references (i.e., cross links), personnel skills required, duration of the task, how long the system is expected to be unavailable while the maintenance is performed, spare parts usage, tools, materials, fluids and lubricants, special test equipment requirements, other miscellaneous requirements, etc., necessary to plan the maintenance activity. Document references and part numbers may be country specific. These tables relationally map much of the knowledge assembled into the TMP documents and distilled from them into the MRCs.

- *Procedures.* These tell human maintainers how to perform scheduled jobs. The procedure document includes text, broken down into tasks and steps; health and safety warnings; cautions and explanatory notes; tables; and graphics detailing job requirements.

The ship-based AMPS system then uses the MRC deliverable to determine when specific maintenance actions need to be performed, flags jobs to the maintenance manager for scheduling, and then assembles the maintenance instruction. The text and required information from the metadata are assembled by AMPS into job-specific printouts for use on the day maintenance is to be performed. Some texts and explicit references are country specific.

Initially, WordPerfect was used to automatically parse and process content from the MRCs to produce electronic deliverables that mapped to AMPS's relational tables and text object stores. WordPerfect helped our authors to capture their maintenance knowledge into a structure they understood, and "single source"⁸ this knowledge into a wide variety of logically different output structures and formats able to be understood by both computer systems and human readers. And we did not have to buy new software to do it (Hall, 2001a).

3.4. *Fundamental issues with word processing*

Although the structured approach using WordPerfect merge tables was hugely successful for one ship, there were major issues that became worse as the volume of repetitive content grew:

- *Proprietary format.* Proprietary word-processing systems are unsuited for maintaining long-lived knowledge.

Word-processing systems become obsolete so rapidly that we needed to reinvent our knowledge repository at least once before delivering the last ship — and the client would have to maintain content for another 27+ years (the programmed life of an ANZAC ship).

- *Delivering massively redundant content into a normalised relational database.* Another fundamental issue became obvious when we started documenting Ship 02 for the RNZN. Although the RAN and RNZN ships were 95 per cent similar at the systems level, some systems differed. We also found it difficult to track engineering configuration changes through ship-specific documents. Even though some metadata was automatically verified, we could not manage configuration information consistently across a growing number of MRCs per ship or similar but not identical routines between ships. Metadata difficulties were highlighted when AMPS rejected records containing inconsistent key data, even though a human maintainer probably would not even see the inconsistency and would have no problems performing the required maintenance.
- *Labour requirements for maintaining massively redundant data.* The final issue with maintaining configuration-specific maintenance knowledge in a flat file word-processing environment is that much of the content is redundant (i.e., content in one document is repeated elsewhere, sometimes hundreds or even thousands of times). With flat files, where a change is required, each instance of the same content must be separately edited. One consideration is the amount of time it takes to manage and effect one change, the times the number of places where the same information needs to be changed. Another is the difficulty to ensure that each redundant change is made consistently.

3.5. *Standards and technologies*

Standards and technologies have been developed to facilitate information storage and transmission in non-proprietary formats. Standard generalised mark-up language (SGML — ISO 8879, 1986) and its derivative XML are particularly relevant for structured documents. SGML (and XML) are languages for defining semantically significant labels for elements of text, and occurrence rules for where each kind of element can be used. A human and computer readable document type definition, or DTD, establishes the rules. The DTD also defines sequential and hierarchical rules for where particular elements may be

⁸"Single source" is a term used by technical writers for applications or processes able to deliver different formats (e.g., on-line help v.s. printed documents) from the same content source. For some, the term includes extracting content from a single source file into qualitatively different output documents as we did here.

allowed or required to occur relative to the overall document structure. Structured applications formally separate content authoring from formatting content for presentation. SGML facilitated the development of content management systems able to semantically identify, process and manage elements of content within documents. HTML, which forms the basis of today's World Wide Web, is defined by a particular SGML DTD, first released to the world by Tim Berners-Lee in 1991 (Sears, 1998, W3C, 2002a). XML, which is currently replacing HTML for encapsulating Web content with a semantically useful markup language, is a somewhat simplified and standardised derivative of SGML (W3C, 2002a).

In 1994/95 Tenix began authoring TRS documents in a stand-alone SGML-authoring environment. This proved the advantages of separating the capture of knowledge as document content from issues of formatting the content for delivery (in a word-processing environment authors of complex documents often spend 20–50 per cent of their keyboard time trying to resolve formatting issues).

4. Document and Content Management Solution (DCMS)

4.1. *Maintaining maintenance documentation is a demanding process*

The client continued to escalate their concerns about problems arising from our use of the “flat-file” word-processing system to deliver data into the relationally based AMPS maintenance management system. They also informally expressed a strong desire for us to standardise the procedural text to facilitate training, by describing common tasks using the same wording, wherever these tasks occurred in different maintenance routines. The 1998 Westralia and Longford disasters focussed even more pressure on the content of maintenance instructions. It was no surprise that Tenix felt the need to individually review some 8000 existing maintenance routines for the four ships to upgrade and standardise all health and safety information prior to delivering the documentation for Ship 05.

Ship 05 delivery, scheduled for March 2001, included a major contract milestone for the full suite of maintenance documentation, which had to be delivered by October 2000 to give the client time to review the documentation before accepting delivery of the ship. We were under notice that if the “data quality” issues relating

to the maintenance documentation package were not resolved, the ship would not be accepted. Contractually, this would have made us liable for major liquidated damages claims for failing to meet the ship delivery milestone.

4.2. *Implementing the document and content management system (DCMS)*

In 1998, Tenix funded a year-long R&D exercise to develop a detailed requirements specification and implementation plan for the selected solution and to evaluate content management solutions for the maintenance documentation. The options considered included migrating to a Microsoft-application-based solution, developing a bespoke database application, or implementing an SGML content management solution. CSIRO's Mathematical and Information Sciences Text Information Management Group, two independent systems implementers with SGML skills, and two suppliers of SGML-based content management systems assisted our evaluation. In the end we selected Aspect Computing to implement RMIT University's Structured Information Manager (SIM), now branded TeraText.⁹ Implementation began in May 1999. Hall (2001a) gives details of the initial solution.

Content in the TeraText DCMS environment is stored in ASCII SGML or XML format and can be downloaded at any time for transfer to other standards-based systems — which resolves the issue of storing and maintaining long-lived knowledge in proprietary word-processing formats.

Procedure text for the MRC is authored in an SGML editor,¹⁰ which requires all work to conform to the semantically structured DTD for that type of document. Metadata authoring can be done either in the authoring environment or within the DCMS browser interface. Authoring metadata in the browser-based TeraText environment is assisted by powerful functions to validate the form of the content and its consistency with master data (e.g., all codes are validated against lists of valid codes for that particular logical element). Sophisticated query facilities also allow authors to rapidly locate potentially reuseable texts to facilitate standardisation.

The native SIM/TeraText product provides highly sophisticated and semantically sensitive tools for indexing and rapidly retrieving content. These indexing and retrieval tools have been built into the specific management environment provided by the DCMS solution; but in

⁹<http://www.teratext.com/>.

¹⁰Tenix currently uses Adobe's FrameMaker + SGML, but we demonstrated that TeraText will work with any SGML/XML editor simply by changing the Windows mime type. TeraText always validates conformance to the DTD on check-in and check-out.

an environment where general content is managed, they offer considerably more sophistication in retrieval and comparable scalability to large volumes of documentation to the retrieval facility provided by the Web search tool Google.⁹

Workflow is fully controlled and tracked in the DCMS environment. New MRCs are started when the supervisor raises a job and assigns it to an author. Authors begin new documents either by cloning and checking out a blank template or (more often) an existing document that is similar to the desired end product. On demand, or whenever the document is checked in or checked out, the DCMS automatically verifies and validates the document's logical structure against the DTD and links, metadata and configuration sensitive information in the documents against appropriate master data repositories. (To speed validation, an overnight automatic extract of the ILSDB is stored within a DCMS repository.) When an author completes a draft, it is checked back in and automatically designated for review and released into an "unallocated" work tray accessed by all authors. Other authors electronically review, annotate the document to highlight any issues, and sign off as having completed their reviews. When sufficient reviews have been completed, the supervisor sets the peer review to complete, which instantly releases the annotated draft back to the original author for corrections. When corrected, the document is released back to the supervisor for a quality check. The supervisor then releases the document to the client for a pre-release review, where the client has the option to annotate any difficulties. Where client changes are required, the supervisor routes the document back to the author for further editing. If the client accepts the document with no changes, the supervisor releases the document for publication.

Compared to the paper-based workflow that the DCMS system replaced, where reviews took weeks or months, and folders were sometimes lost or misplaced and the process had to start over, the electronic workflow is immediate, and any delays are immediately traceable via reports and queries. Several reviewers can work simultaneously on the same document, and the speed and traceability of the process ensure that the review processes are completed to a very high quality — in less than a day if required.

4.3. *Managed links and annotations as tools for converting tacit knowledge to explicit*

DCMS was updated in 2001 to add tools for capturing contextual knowledge (Fig. 5). Links and annotations associated with specific elements of information in doc-

uments relate these elements to other information that turns information and tacit knowledge into assimilated explicit knowledge. DCMS contains two specialised repositories and a registry to facilitate this assimilation process.

- *Source repository.* An object store able to accept and preserve any binary files available in electronic format. Stored files may be launched for viewing in an appropriate external tool.
- *Source registry.* Captures metadata relating to any source document, whether or not the document is available in electronic format and preserved in the source repository. In addition to providing standard information about document sources and version numbers, comments allow qualifications to be recorded against the source documents. New versions of a document can easily be identified by changes to the standard metadata.
- *Content repository.* Repository for deliverables and a special class of XML records called annotations. Each annotation minimally links to a content element in a deliverable document and includes metadata describing circumstances of the annotation's creation, free text for the author's comments, and optional link(s) to deliverables or source registry item(s).

There are three kinds of annotations:

- *Author.* Used by authors to describe the sources of information underlying the assimilated content of a text element. If sources are documents, the author is expected to register them in the source registry, and if available electronically, capture the document itself into the source repository. Authors are instructed to note what specific information from a source is used in drafting a deliverable (i.e., chapter and section or page number).
- *Reviewer.* Reviewer comments intended to be addressed at the rework stage by the document author.
- *Client.* Comments made by the external client, for attention of the document supervisor and author.

Author annotations remain permanently associated with the document when it is released for publication (although they are stripped from the document on delivery to the client). Reviewer and client annotations remain associated with the particular archived version of the document for audit and change management purposes, but are disconnected from the published version of the document so that each new round of revision begins with a clean slate. Annotation capabilities in the DCMS have important implications for knowledge management because they provide a unique tool to convert, capture and codify implicit knowledge to make it explicit.

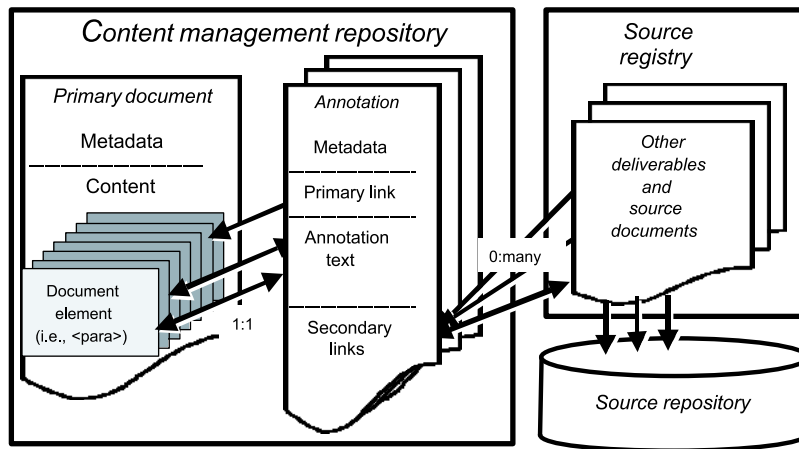


Fig. 5. The architecture of an annotation.

The author’s implicit knowledge of the original information sources is a highly valuable component of deliverable documents. Yet such information is easily lost with staff movements unless means are provided to capture the knowledge as the documents are drafted. The ordinary paper or word processor-based technical or commercial authoring environment provides no practical capability to capture such author knowledge other than through explicit references in the text. The value of such implicit knowledge becomes obvious when a new version of a source document is received from an equipment manufacturer (e.g., the gas turbine engine), possibly years after the deliverable document was drafted. Only the original author has any World 2 (i.e., personal) knowledge as to whether the changed source document is likely to require changes to particular deliverable documents. Anyone else may have to spend days reading a range of documents in the hope of determining which ones are likely to be affected, and the chances are that critical impacts will be missed through unfamiliarity with the documents being checked.

With DCMS, checking source registry metadata quickly reveals whether the newly received document represents a revision of something previously referenced. A “where used” report on the originally registered document displays or prints out author annotations referencing the older version. Annotations will identify what information in the source was used (hopefully to page number). It is then very easy for anyone with even a moderate knowledge of the system to which the document pertains to follow the link from the annotation to the element in the deliverable that used the information and back to the particular information in the source document used, to determine whether the source document has changed in a way that requires the deliverable to be changed. This

changes a fallible “impact analysis” task that could take weeks to complete when done manually, into a largely automated one that can be done to a much higher degree of reliability in minutes.

Annotations can also be used capture and reference undocumented contextual drivers (e.g., meetings, telephone conversations, etc.) for document changes, which can be invaluable in resolving disputes and for recording lessons learnt. Annotations turn implicit knowledge of contexts and circumstances into indexable, searchable and easily retrievable explicit knowledge.

5. BENEFITS in the Logistics World

5.1. Metrics

As noted above, Tenix faced a major challenge to acceptance of our ILS deliverables if we failed to resolve outstanding problems with the MRCs. Our experience with word processing and contractors was that, irrespective of how many contract authors we hired, we could not have reworked 8000 individual documents to “standardise” health and safety texts (to say nothing of other kinds of texts) into a deliverable product in the time available. Temporary staff had their own ideas about how documents should be structured and inevitably did things that caused parsing failures in the WordPerfect merge/macro delivery.

Converting content from WordPerfect into SGML and delivering it to the client proved the technology’s value in managing World 3 knowledge (Hall, 2001a). DCMS eliminated much of the redundant text. Because most MRCs were identical across most ships (except for configuration details and “language” differences), we were able to collapse over 8000 ship-specific documents for

four ships into about 2000 master documents covering ten ships for two navies. “Single source” techniques managed configuration and language variants within a single master document. In our best case, 46 separate equipment and ship-specific MRCs covering four ships collapsed into a single master MRC covering 10 ships. Master documents were resolved into sets of ship-specific documents by the DCMS delivery process. Automatic validation in authoring and delivery processing also guaranteed the consistency of the configuration sensitive information and its validity against the ILSDB.

The complete process to produce ~2000 edited, valid and signed-off deliverables from ~8000 “raw” SGML files took five authors approximately 3000 person hours. As author skill grew in the conversion task, the average hands-on processing time (including all stages in the workflow) per document reduced to around half an hour. Document content was displayed in a logical structure view as well as formatted text, and we had a query facility to call up in seconds all related documents for side-by-side comparison. With such tools, we were able to very thoroughly proof, edit and in some cases even completely rewrite the routines to standardise document structures, warnings, cautions, notes and other commonly used texts so that they were consistent across documents.

Compared to our sophisticated word processor, apart from reducing hands-on processing time per document by much more than 50 per cent, we reduced the number of routines we had to deliver for Ship 5 from an anticipated 10,000 to around 2000. With word processing, the only way we came close to delivering a consistent set of documents was to deliver a complete set to each ship every year. With DCMS we now deliver only documents that change, and we do this so that the changed documents become effective in complete synchrony with the completion of a related engineering change. Aside from meeting applicability and effectivity requirements, we reduced the cycle times to implement document changes from a year to a couple of days (if required), and reduced subsequent data deliveries by more than 95 per cent. Apart from the magnitude of improvement in production metrics, separating explicit knowledge from delivery formats gave us similarly massive (though difficult to measure) improvements in quality.

5.2. Closing the knowledge circle to better manage fleets and facilities in operation

Managing knowledge in a virtual World 3 paradigm also has major implications for project operators. Fig. 6

summarises how fleet maintenance knowledge is currently managed at the fleet level, using additional Tenix-developed software derived from the OARRS application that we used to meet our ILS TE&V requirements. Most of the management processes illustrated take place more or less automatically with minimal human involvement.

Tenix’s logistics analysts and technical authors continue to maintain maintenance routines, acting as agents for the ANZAC Ship Alliance, a consortium comprising (a) Tenix — responsible for hull and mechanical systems, (b) SAAB Systems — responsible for many of the combat and electronic systems, and (c) the ANZAC fleet System Program Office — responsible for in-service support of the Australian fleet. While ships are still being built, Tenix maintains master configuration data (i.e., details of specific items of equipment for each ship).

Maintenance routines with their associated metadata are delivered electronically into a central AMPS system, and the applicable content is allocated to the respective ships to take effect automatically along with any engineering changes that become effective. AMPS knows the engineering configuration of the ship to know which versions of the recorded knowledge relate to which equipment on each ship, and it automatically monitors the calendar and a variety of ship operational variables (e.g., running hours, oil pressures, etc. relating to particular equipment) to flag which maintenance routines are due to be performed (i.e., “triggered”).

The knowledge as to when the particular maintenance should be performed in relation to the variables being monitored is translated into World 3 knowledge and logic by the technical author drafting the MRC, while the decision based on this knowledge as to when a particular maintenance activity is to be triggered is made automatically by the AMPS system.

The onboard maintenance manager (a human) then assigns flagged jobs to maintainers to be performed on a certain day. AMPS prints out configuration-specific details of the job, which it assembles from several disparate sources in its repository (many derived from MRC content, some from Navy sources): what is to be maintained; what repair parts and other materials are to be taken to the job; what is required in terms of tools, test and miscellaneous equipment; the step-by-step details of the procedure to be performed; and any job completion forms that are required to capture measurements or other observations on the job. The maintainer(s) complete the job and return the completion report for entry back into AMPS. This includes details of any anomalies. In the case of a failure (i.e., unscheduled maintenance), a job is also raised in AMPS that includes completion of a report describing the

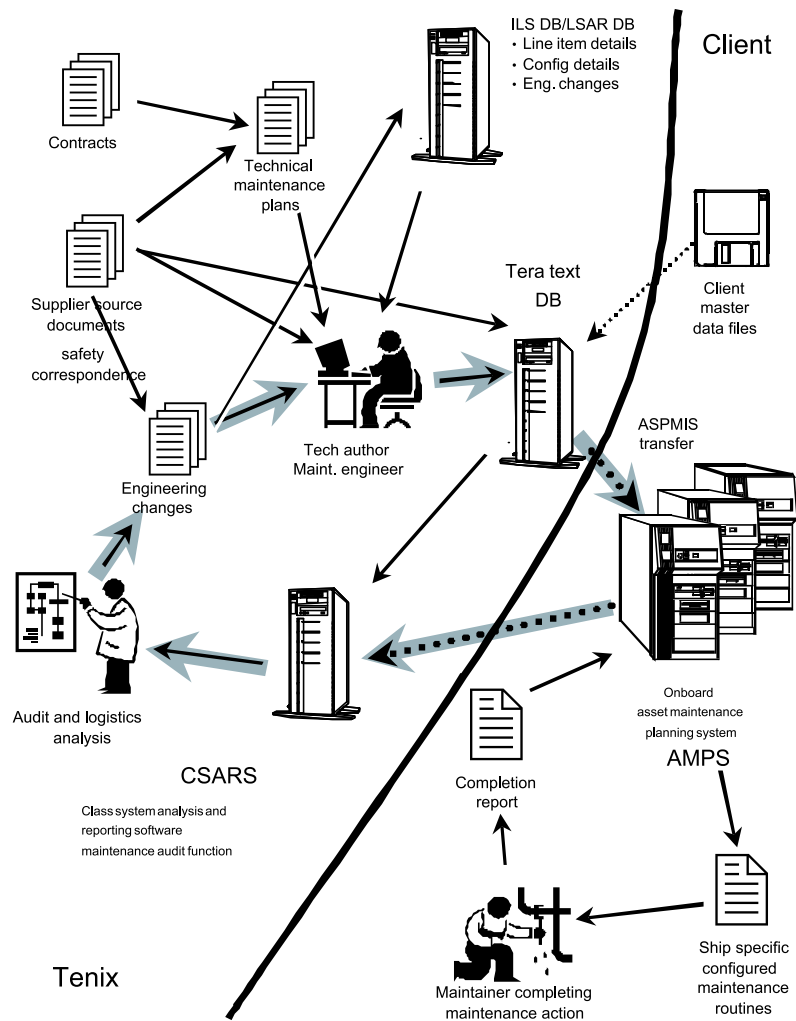


Fig. 6. Closing the circle. Implementing corrective feedback into the fleet knowledge management cycle.

nature of the failure, what was done to correct it, and how long the system was out of service.

Operational data is automatically extracted from AMPS into a Tenix-developed tool called CSARS (class system analysis and reporting software) based on the OARRS technology which Tenix developed (Shelley, 1996) to determine whether ANZAC ship systems were meeting their contractual availability thresholds. CSARS is used to measure and analyse system and equipment performance in service. The AMPS system captures the raw information in the form of job results, time to complete, spares used and other recorded parameters. CSARS processes the AMPS observations along with data from other sources to calculate performance measures for the logistic support knowledge and highlights systems that appear to be performing poorly for human attention. Logistics analysts then use CSARS and other analytical tools to drill down to examine parameters recorded by specific jobs

either at the equipment level where failures occur or their impact on system and ship availability as determined using CSARS's availability hierarchy logic (Shelley, 1996). Logistics analysts then use this largely automated analysis of operational performance to feed knowledge back into the logistic support system to optimise fleet performance on the basis of cost/benefit trade-offs. Support adjustments may include reengineering parts that fail unexpectedly frequently, changing the number and distribution of particular repair parts in the support supply chain, changing the frequency of certain kinds of scheduled maintenance, adding new maintenance procedures or changing details of existing maintenance procedures. The result of implementing this feedback loop is that operating and maintenance knowledge is continually tested and validated against the performance of the ships in the real world in such a way that the operating costs to maintain the ships to a required degree of availability can be

optimised in something close to real time, and the resulting knowledge is also preserved and available for incorporation in new projects as these are put out to tender.

6. Conclusion

As demonstrated in the case history, large engineering projects generate large volumes of knowledge that must persist and be securely managed for much longer than the expected involvement of any particular individuals in the project. The nature of this knowledge and the requirements to manage it can best be understood in a framework of Karl Popper's three worlds of knowledge rather than in any framework based solely on Michael Polanyi's personal knowledge that encompasses only some aspects of Popper's World 2. Knowledge required to operate and maintain the assets produced by the project may have life and death implications for the operators of the engineered product, as sadly demonstrated in the Longford gas plant and the Westralia supply ship disasters. Such operational and maintenance knowledge is also of great economic value for the engineering product's longevity in service and life-cycle costs. World 2 personal knowledge may suffice for safe day-to-day operations, but it is not humanly possible to retain as personal knowledge the knowledge that is required to deal with all possible unusual events. However, it is important to provide systems that can locate and rapidly recall critical knowledge from World 3. Thus, as far as possible, the overall knowledge production and management strategy for the long-lived engineering product must provide the means to test and continually monitor the performance of the knowledge-based documentation products against the actual performance of the product in service in the real World 1. Costs and schedules to produce and maintain the knowledge are also of major economic concern to the engineering organisation responsible for delivering the product and its associated knowledge base — not only must the knowledge suffice for the client's operational requirements, but the supplier must be able to make an economic return on the costs to produce the knowledge. Thus, the processes and technologies used to capture, manage, deliver and retrieve knowledge relating to the operation and maintenance of engineering products are of major concern to both suppliers and operators of the products.

New technologies implemented by the Navy (e.g., AMPS) and Tenix (TeraText and CSARS) have changed the way fleet maintenance knowledge is managed (i.e., replacing ponderous, linearly organised paper manuals that are expensive to produce and difficult to update with virtual content that can be processed, distributed and retrieved at high speed) to substantially improve the

quality of knowledge and safety information in documents and substantially decrease the cycle times and costs to produce and maintain the knowledge. Realising the savings for new projects will allow clients to spend more of a limited budget on capabilities and better operate existing capabilities.

However, maintenance knowledge is only half the story. The development of computerised aids to cognition, by replacing paper with electronic technology able to do smart things with contextually relevant content (i.e., knowledge), independent of its active presence in human brains or its representation on physical paper, is revolutionising the way organisations like Tenix conduct their business.

Tenix is currently implementing structured authoring and content management applications in the earlier stages of the project cycle to manage tender and contract documentation. Here, we hope to work in collaboration with the Australian Defence Materiel Organisation people to develop XML standards for exchanging bidding and contract documents that will facilitate implementing content-based knowledge management applications on both sides of the contractual interface. Corporate management is only starting to understand what has been done in the logistics area. However, our experience with maintenance knowledge suggests that we can expect comparable improvements in the quality of the knowledge encapsulated in bids and contracts, and achieve similar reductions in costs and cycle times to produce and deliver these documents as we have achieved with maintenance documents. The kind of rapidly adapting, agile organisation we hope to become through the application of technologies and processes to work with and use more effectively the vast amount of World 3 knowledge our company works with was given the name "cybercorp" by James Martin (1996).

Effective knowledge management maximises organisational efficiency, and those who do it better and faster will win competitions at the expense of those who do not. Systems like TeraText and AMPS provide new cognitive tools for query and discovery (*know what you know*), structured and controlled authoring and processing of elements of content (*eliminate redundancy to reuse what you know*), electronic routing for review, approval and delivery to the end user (*better QA and progress at high speed*), instant change management tracking and reporting (*know where and why you did it*), and contextual links to source data (*converts implicit often lost knowledge to explicit and easily retrievable*). This applies both to suppliers like Tenix, who can complete their business cycles faster and with less labour, and to the naval operators who can anticipate faster delivery cycles and who

will have the knowledge available to help them maximise the availability and effectivity of their fleet assets when most required and do this within always constrained budgets.

References

- Blanchard, B. S. (1992). *Logistics engineering and management* (4th ed). Englewood Cliffs, NJ: Prentice Hall.
- Capurro, R. (2002). Skeptical knowledge management. In H.-C. Hobohm (Ed.), *Reader: Knowledge management and libraries*. Munich: IFLA (International Federation of Library Associations and Institutions) Publication series. Saur. www.capurro.de/skepsis.html.
- Clarke, M. (2000). *Longford Royal Commission into the explosion and fire on 25 September 1998 at the Esso Gas Processing Plant*. 2000 Queensland Mining Industry Health and Safety Conference, 27–30 August 2000, Townsville, Queensland, Australia. http://www.qmc.com.au/docs/conferences/QMC_2000/conf.clarke.pdf [Valid 7/10/2002].
- Commonwealth of Australia (1998). *Report of the Board of Inquiry into the fire in HMAS Westralia on 5 May 1998*. Canberra ACT 2601: Director of Coordination — Navy, R1-4-C070, Russell Offices. <http://www.navy.gov.au/fleet/O195westralia/boi/report.htm> [Valid 12/08/2002].
- DAU (undated). Interactive Electronic Technical Manuals. Defense Acquisition University. US Department of Defence. <http://log.dau.mil/ietm-toc.asp>.
- Dawson, D., & Brooks, B. (1999). *The Esso Longford gas plant accident: Report of the Longford Royal Commission*. Melbourne: Government Printer for the State of Victoria. (<http://www.bookshop.vic.gov.au/infovic/anonymous/products.asp?ProductID=24419>). For a summary see <http://www.ourcivilisation.com/decline/gasbang.htm>.
- DNV (2000). DNV leads development of process industry safety. DNV Forum, 2000(2). http://www.dnv.com/dnvframework/forum/articles/forum_2000_02_12.htm.
- Firestone, J. M. (2000). *Knowledge management: A framework for analysis and measurement*. White Paper 17. Executive Information Systems, Inc. <http://www.dkms.com/KMFAMrev1.PDF> (65 pp.).
- Firestone, J. M. (2001). Key issues in knowledge management. *Knowledge and Innovation — Journal of the Knowledge Management Consortium International, Inc.*, 1, 3. <http://www.dkms.com/FirestoneissuesKIv1n3.pdf>
- Gaines, B. R. (2003). Organizational knowledge acquisition. In C. W. Holsapple (Ed.), *Handbook on knowledge management: 1. Knowledge matters* (pp. 317–347). Berlin: Springer. <http://www.repgrid.com/reports/KM/OKA/>.
- Hall, W. P. (1983). Modes of speciation and evolution in the sceloporine iguanid lizards. I. Epistemology of the comparative approach and introduction to the problem. In A. G. J. Rhodin, & K. Miyata (Eds.), *Advances in Herpetology and Evolutionary Biology — Essays in Honor of Ernest E. Williams* (pp. 643–679). Cambridge MA: Museum of Comparative Zoology.
- Hall, W. P. (2001a). *Maintenance procedures for a class of warships: Structured authoring and content management*. Technical Communication 38:3. pp. 235–247. <http://www.tenix.com/PDFLibrary/91.pdf> [Valid 8/08/2002].
- Hall, W. P. (2001b). Content management in a defence prime contracting environment: Present and future. Open Publish 2001 — Sydney, 31/7–2/8, 2001. <http://www.binarything.com/binarything/openpublish/OpenPublish2001a.pdf> [Valid 8/08/2002].
- ISO (1986). ISO 8879 Information processing — Text and office systems — Standard generalized markup language (SGML). International Standards Organization [see <http://www.standards.com.au/catalogue/script/Details.asp?DocN=stds000004623> Valid 11/08/2002].
- Kuhn, T. S. (1962). *The structure of scientific revolutions* (173 pp). Chicago: University of Chicago Press.
- Kuhn, T. S. (1970). Postscript — 1969. In *The structure of scientific revolutions* (2nd ed.). Enlarged (210 pp). Chicago: University of Chicago Press.
- Kuhn, T. S. (1977). Second thoughts on paradigms. In Suppe (Ed.), *The structure of scientific theories* (2nd ed.) (pp. 459–482). Chicago: University of Chicago Press.
- Kuhn, T. S. (1983). Commensurability, comparability, communicability. In *PSA 1982: Proceedings of the 1982 Biennial Meeting of the Philosophy of Science Association, Vol. 2*. From Kuhn (2000). The road since structure. Chicago: University of Chicago Press. Longford Royal Commission (1999).
- Martin, J. (1996). *Cybercorp: The new business revolution* (326 pp). Amacom.
- Moss, M. W. (2002). *Why management theory needs popper: The relevance of falsification*. Paper presented at the Developing Philosophy of Management Conference, 26–29 June 2002, organised by Reason in Practice, The Journal of Philosophy of Management and the Forum for European Philosophy.
- Nickols, F. (2000). The knowledge in knowledge management (KM). In J. W. Cortada & J. A. Woods (Eds.), *The knowledge management yearbook 2001–2002*. Butterworth-Heinemann. http://home.att.net/~nickols/Knowledge_in_KM.htm [Valid 8/08/2002].
- Nonaka, I. (1996). Knowledge has to do with truth, goodness, and beauty. In *Dialog on leadership: Conversation with Professor Ikujiro Nonaka, Tokyo, Japan, February 23, 1996 and with Claus Otto Scharmer*. <http://www.dialogonleadership.org/Nonaka-1996.html>.

- Nonaka, I. & Takeuchi, H. (1995). *The knowledge creating company: How Japanese companies create the dynasties of innovation*. New York: Oxford University Press.
- Polanyi, M. (1958). *Personal knowledge: Towards a post-critical philosophy* (428 pp). [Corrected Ed., 1962]. Chicago: University of Chicago Press.
- Polanyi, M. (1966). *The tacit dimension*. Routledge & Kegan Paul.
- Popper, K. R. (1934, 1959). *The logic of scientific discovery, sixth impression* (Revised) [First English ed., Hutchinson, 1959. First published as *Logik Der Forschung* in Vienna: Springer, 1934] (480 pp). London: Hutchinson & Co., Ltd.
- Popper, K. R. (1963, 1972a). *Conjectures and refutations: The growth of scientific knowledge* (4th ed.) (Revised) [1st ed., 1963] (431 pp). London: Routledge and Kegan Paul.
- Popper, K. (1972b). *Objective knowledge: An evolutionary approach* (380 pp). London: Oxford University Press.
- Saint-Onge, H. (1996). Tacit knowledge: The key to the strategic alignment of intellectual capital. *Strategy & Leadership*, 24, 2. <http://www.knowinc.com/saint-onge/library/tacitknow.htm>.
- Sears, P. (1998). Chapter 13. HTML origins, owners, good Practices. In M. Abrams (Ed.), *World Wide Web: Beyond the basics* (483 pp). Englewoodcliffs, NJ: Prentice Hall. <http://ei.cs.vt.edu/~wwwbtb/book/chap13/>.
- Shelley, T. (1996). *Data management in support of measuring operational availability of complex systems*. MSc Minor thesis, School of Communications and Informatics, Victoria University of Technology.
- Sheppard, N. (1999). Michael Polanyi and the philosophy of science: The viewpoint of a practising scientist. *Polanyiana*, 8, 1–2 [Originally published in *Appraisal*, 2(3) (March 1999), 107–115]. <http://www.kfki.hu/chemonet/polanyi/9912/sheppard.html>.
- Snowden, D. (2002). Complex acts of knowing: Paradox and descriptive self-awareness. *Journal of Knowledge Management*, 6(2), 100–111.
- Sveiby, K.-E. (1994). *Towards a knowledge perspective on organisation*. Doctoral dissertation, University of Stockholm, S-106 91 Stockholm. <http://www.sveiby.com/articles/Towards.htm> [Valid 8/08/2002].
- Sveiby, K.-E. (1997). Tacit knowledge — An introduction to Michael Polanyi. In Sveiby Knowledge Associates. <http://www.sveiby.com/articles/Polanyi.html> [Valid 8/08/2002].
- Sveiby, K.-E. (2000). Frequently asked questions. In Sveiby Knowledge Associates Web site — <http://sveiby.konverge.com/articles/FAQ.htm> [Valid 8/08/2002].
- U.S. Department of Energy Office of Science (2003). Human Genome Project Information. http://www.ornl.gov/TechResources/Human_Genome/W3C (2002a). Extensible markup language (XML). W3C Architecture Domain. World Wide Web Consortium. <http://www.w3.org/XML/> [Valid 11/08/2002].
- Wilson, T. D. (2002). The nonsense of “knowledge management”. *Information Research*, 8(1). <http://informationr.net/ir/8-1/paper144.html>.

William Hall is Documentation Systems Analyst for Tenix Defence and an Honorary Research Fellow in Monash University’s School of Information Management and Systems.

Dr. Hall earned his PhD in Evolutionary Biology from Harvard University in 1973 and has held teaching and research positions in biology through 1980. From 1977–1979, he spent much of a postdoc in genetics at the University of Melbourne researching the history and epistemology of evolutionary biology. From 1981 Dr Hall became involved with technical writing and documentation management in the computer industry. Since 1990 he has had a number of roles in Tenix Defence and its progenitors relating to document content management.