cih Solutions
SERVICE MANAGEMENT INTEGRATORS

Knowledge Management within ITSM

The first in a series of white papers from CIH Solutions that discuss topical issues in IT Service Management

Mike Simpson
info@cihs.co.uk
June 2015

Abstract

This white paper discusses how Knowledge Management (KM) can be used to manage risk and control costs in an IT Service Management environment.

This white paper discusses how Knowledge Management (KM) can be used to manage risk and control costs in an IT Service Management environment. The paper identifies four 'hot spots' based on the author's experience and outlines common problems and suggests solutions using KM.

Introduction

As with most terms found in IT the term Knowledge Management means different things to different people. There is much available on the subject of KM and the term is often interchangeable with other terms such as intellectual capital, information management, data management and document management. In reality, KM embraces all of these.

So, what is my definition of KM in relation to an ITSM organisation?

First, this is not about scale. A KM system can operate just as effectively in a small organisation as a large enterprise. The principles remain the same – identifying, collating, storing and retrieving knowledge for use by all personnel in their day-to-day tasks. Also, this is not just about documents and data. When the experience of personnel is added into the mix we get Knowledge and this needs to be captured and stored for future use.
Second, from my experience the key feature of a KM system within an ITSM organisation is the understanding that different information has different values depending on circumstances. For me assigning value to information is vital and has priority over the capture of all available material.

At this point I should add that I do not differentiate between an MSP serving external clients and an internal IT service provider. The same KM principles apply. Also, the KM system described in this paper should be considered a 'practical solution' that can be implemented with limited resources and budget and extended over time.

I want to begin by briefly describing two KM systems that I have encountered in the course of my consultancy work.

Example One

I've seen only one truly outstanding example of an enterprise wide KM system and that was at a European pharmaceutical company. What struck me about this KM system was the sheer scale of the repository containing research papers, trials results and project documents covering decades of research amounting to many millions of pages and database entries. The success of this KM system was of course the strength of the underlying thesaurus that enabled scientists to discover (or perhaps re-discover) knowledge to support the design of a new R&D programme.

Example Two

My second example is at the other end of the scale. This is a local KM system that supports an IT organisation that provides hosting support for external SAP clients. This KM system also impressed me but for a different reason. Without any real top down sponsorship or funding the technical teams had created their own KM system based on a single central repository, but where all the content was created, published and maintained under very strict guidelines by a few key members of staff, but accessed by many. The rationale for using this approach was to bring discipline to the management of documents and data that were considered vital to the successful running of their IT organisation.

KM Model for ITSM

The rationale for the second example above sounds somewhat obvious, but the background problem as explained to me was one of long term ill-discipline in the day-to-day management of key information.

Individuals, both staff and sub-contractors, would create multiple documents, held in isolated repositories or held on local drives, resulting in poor retrieval and inaccurate information.
The problem is a familiar one. Admittedly, this KM system is basically document management, plus some other information formats and a simple data classification system, but in my view this doesn't matter as the problem of badly managed information was controlled by introducing a strong KM framework with a central repository to address a specific local need.

It is this model of KM that I want to discuss as the starting point for KM for ITSM, but first I need to say something about the concept of assigning value to information.

Defining Business Value

I mentioned above that assigning value to information is vital.

> "Whilst all information is valuable, depending on circumstances, some information suddenly becomes more valuable".

I call this category High Business Value information. So, what does it mean exactly? Essentially, this is a category of business information that covers all the vital and irreplaceable business records, documents, information and data that are associated with sensitive areas like customer data, compliance, security, personnel, finance and legal and commercial activities.

It is this category that has the potential to damage an ITSM organisation should this material be compromised by loss, breach of security, inaccuracy or the inability to locate and retrieve quickly when needed. It is the failure to identify, capture, publish and retrieve this category of knowledge that can have a significant impact on the management of risk and cost control.

Whilst all information is valuable, depending on circumstances, some information suddenly becomes more valuable.

KM Framework

Our first step is to build a KM Framework. This framework must define the KM life cycle to create, capture, review, release, amend, publish and retire content. In addition, the KM Framework must define a system of classification for the ITSM information. We have already identified a need to segregate high value information – I'm calling this Layer 1 information. All the remainder of the ITSM information and data is collected into Layer 2.

Basically, for Layer 1 we know what we want and where it is – hence we can find it quickly using a hierarchy with a controlled vocabulary where everything is tagged.

However, for Layer 2 the structure is more linear using a Thesaurus and non-controlled vocabulary. This allows for a more 'search and discover' approach.

Finally, the framework will identify the ITSM knowledge managers who will be responsible for implementing the framework, plus a KM Steering Committee.

Five Stages of the KM Framework

There are five stages within the KM Framework and these are shown in Figure 1 below. By following this five stage sequence all the information considered as High Business Value can be identified and either uploaded into the KM Database or retained in local repositories (known as source databases). This is the Integrate stage that is covered in detail later on under the Hot Spot scenarios.

Each stage should be followed for Layer 1 and then repeated for Layer 2.

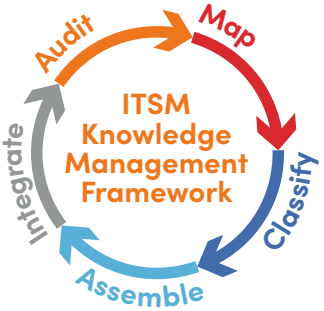

Figure 1 – Five Stages of KM Framework

Audit – once the categories within Layer 1 have been identified all the material to be included in each category needs to be identified. The audit will do this and will cover different media formats such as PDF, database tables, e-mails, webinars and HTML et al.

Map – during the audit the location of the material is identified. This is mapping and will be needed when the KM database is designed and built to identify what material should be transferred to the KMDB and what material should remain in local repositories.

Classify – once all the information has been identified for the categories of Layer 1, the documents and data can be classified according to the controlled vocabulary system and the hierarchy structure.

Assemble – once classified and physically located, the content for each category should be assembled as a schedule of descriptive metadata tables complete with information titles, document numbers, versions, data sets and physical location.

Integrate – once all the information has been assembled the metadata tables can be used to manage the population of the KMDB – either directly with content or connected to other repositories to extract the content. These are known as source databases.

Classification

As mentioned above it is important to classify by value as well as classify by subject. For example, all customer data should always be considered high value, but the exact list will depend on the types of client and services that are supported by the ITSM organisation.

When it comes to the subject of classification there are many standards1 on taxonomy and debates about linear versus the hierarchy structure approach. I'm therefore suggesting that it makes sense to divide our total ITSM information into two distinct groups – the High Business Value information already discussed and a second group which is essentially everything else. I'm calling the first grouping Layer 1 and the second grouping Layer 2.

Once all the information has been divided into these two layers we must structure the information in two different ways. Figure 2 below shows this division.

Layer 1 should be structured using a taxonomy with a hierarchy and controlled vocabulary. This scheme will identify the information according to importance, sensitivity and security level, and will be used to control access to the information in Layer 1.

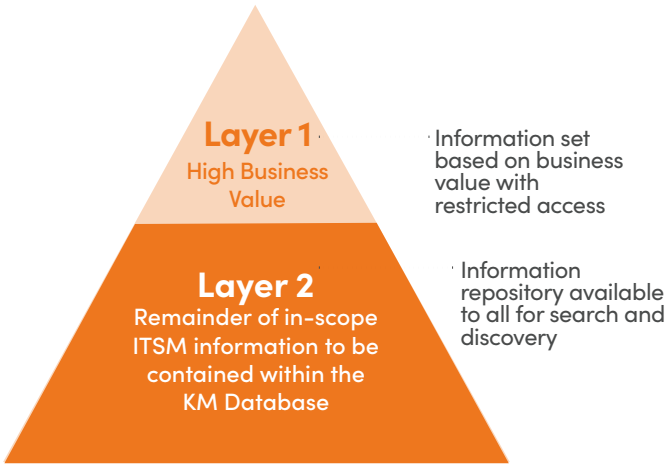

Figure 2 – Grouping Information by Layers

The search tools that underpin our KM system will then be able to locate and retrieve any of the information in Layer 1 very quickly. Layer 1 will typically have the lowest volume.

For our second layer – Layer 2 - I suggest a thesaurus with a more linear structure that will allow more of a free form of search and retrieval based on a smaller number of the terms.

Not everything needs to be tagged in Layer 2, instead broader searches and cross searches can be adopted to allow a more 'search and discovery' approach even 'looking inside' some documents and files to locate content of interest.

This makes sense as the population of Layer 2 will cover all manner of archived project material, design documentation, presentations, non-critical business records et al. Layer 2 will typically have the highest volume.

Hierarchy of Layer 1

Given the relatively simple structure of our KM system I suggest a top down approach for Layer 1, based on a hierarchy of Categories and Sub-categories using a controlled vocabulary to tag documents and data sets. An example is shown in Figure 3 below. As Layer 1 is the primary focus of our initial KM design and build it's not my intention to outline the structure of Layer 2.

Once all the constituents of Layer 1 have been identified during our Audit stage all the information and data can be divided into Categories. These categories will be assembled under various functional headings, for example:

Category 1 – Customer Data
Category 2 – Compliance
Category 3 – Legal
Category 4 – Service Continuity
Category 5 – Finance
Category n – Security

Once all the Categories have been identified then the material should be further sub-divided into Sub- categories. I would suggest that these three drill-downs are sufficient to hold all the information in Layer 1. The Sub-categories will contain all the specific document and data sets that relate to a particular Category and this can be assigned by client or customer type or by any other specific grouping.

This hierarchy is not meant to be in any way prescriptive, just examples on the concept of Categories and Sub-categories.

Example 'Hot Spots'

I've identified four possible 'hot spots' based on personal observations of real life events and these are shown in Figure 4. Clearly, there will be others depending on the set-up of a particular ITSM organisation and the types of client it supports.

The figure is based on a simplified ITSM organisation that could be either a MSP dedicated to external clients, or an ITSM organisation providing IT services to an internal client. The IT Operations can be either internal or external hosting with or without applications support. For the purpose of this paper it is assumed that the IT Operations is in-house and provides hosting, communications and applications support - within an overall governance framework.

Figure 3 – Classification Hierarchy

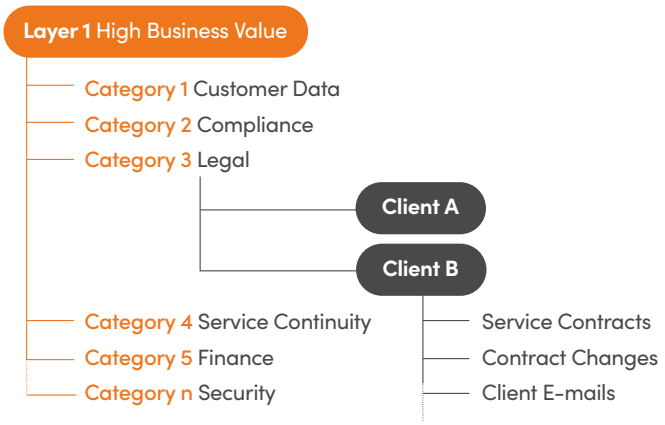

Figure 4 – Example Hot Spots

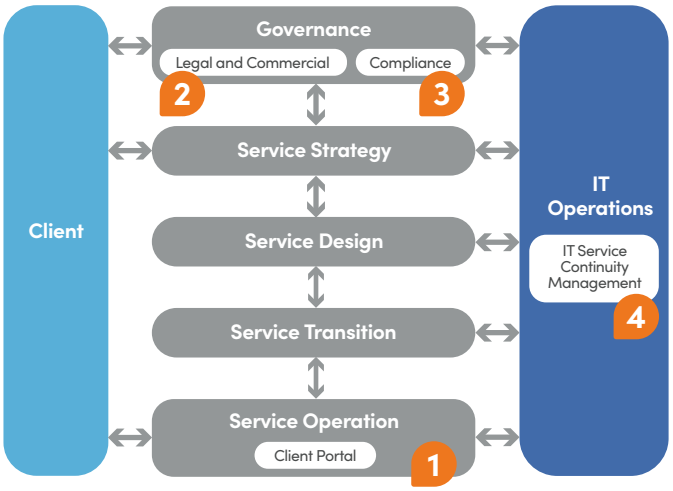

There are four example 'hot 'spots' shown in Figure 4.

- Client Portal – Risk to reputation due to poor quality of customer information
- Legal and Commercial – Cost of litigation due to incomplete contract audit trail
- Compliance – Cost of compliance due to audit failure and forced re-work
- Service Continuity – Risk to IT service continuity due to inadequate preparation

All of the above examples relate to the absence, inaccuracy or timely retrieval of information.

Risk to Reputation (Hot Spot 1)

In this scenario I've created a simple Service Operation (SO) organisation that has responsibility for managing the information available to customers via a Client Portal. I should state at this point that not all of the information available through the portal is the responsibility of the SO team. Some material will be supplied direct from the Client for uploading onto the portal – material from the Marketing Department such as prospectus and application forms.

The remainder of material will be service and technical support information produced within SO and cover such topics as service availability status, technical self-help and how-to-do-it video clips. The client portal also has a secure area for the client customer groups to access data on performance against SLAs.

The 'Risk' we are trying to mitigate here is out-of-date, missing and inaccurate information being posted to the client portal. The current arrangement within our SO is that information is currently held in separate repositories. Information is identified and collected and then manually or semi-automatically uploaded onto the Client Portal database using scripts. The risk here is that:

- not all information is collected at the right time (like monthly SLA data updates)
- incorrect information is selected for the right location
- correct information is uploaded to the wrong location
- not all information is collected

All the above risks can be minimised by the correct processes and checks in place and rigorously enforced. However, experience has shown that this manual and semi-automatic process can break down over time and quality – and reputation – can be impacted.

All the client information that was previously managed manually has now been compiled into metadata tables from the Audit – Map – Classify – Assemble stages. We can now move to the Integrate stage. The metadata tables will hold the locations of all the information and data needed to be accessed by the client portal and the KMDB will use distributed queries to collect all the information and data from these locations. In practice these will be permitted areas within local repositories (or tool set databases) – known as source databases. See Figure 5.

For example, the Known Error database (KEDB) could supply diagnostic help and work-arounds for self- service customers for the most common errors. The KEDB will also collect Event and Incident Management data in support of the SLA reporting that is provided to the client business units via the portal. The Configuration Management database (CMDB) will also be another source database for the supply of data to the client on service configuration.

Cost of Litigation (Hot Spot 2)

My second scenario relates to the threat of litigation as a result of a breach of contract. Whilst this sounds dramatic it is important not to underestimate the legal and commercial requirements to hold and maintain all contractual material and associated business records.

Most service based agreements come with some form of service credit arrangement. However, a decrease in payment may not fully compensate a client for poor service particularly when a number of service failures occur in quick succession or a major outage lasting several days hits not just the client but the client's customers. Such a scenario could be considered a breach of contract resulting in litigation to seek damages and a termination of the service contract.

Figure 5 – KM Integration of Client Portal Information

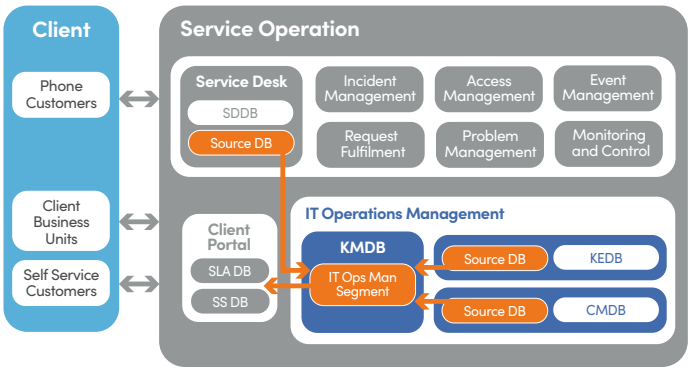

Figure 6 – KM Integration of Legal Information

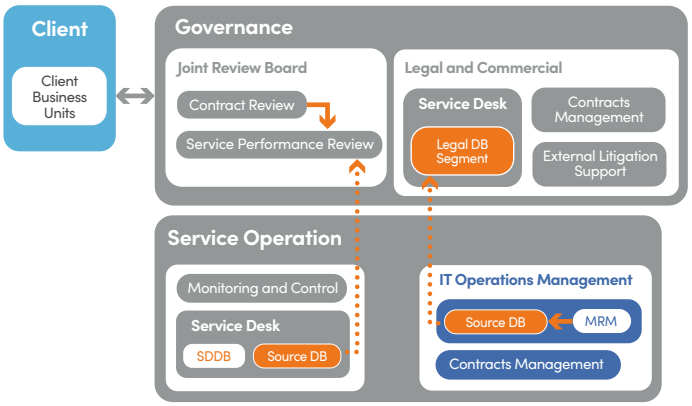

Any move to litigation will result in a demand from the client's legal team for all relevant information to be handed over. This is known as e-discovery2 and the Service Operation team along with the organisation's legal department will need to respond very quickly in a short time frame.

This is another example of how the KMDB can be used to store high business value information. Figure 6 shows how the KMDB can contain a Legal DB segment that is used to store in one location all contractual and historical SLA information relating to an individual client. As with Scenario 1, the metadata tables will hold the locations of all the information and data needed to be accessed by the Legal KMDB segment.

Again, distributed queries are used to collect all the information and data from these source DB locations.

The information will include all versions of contracts, contract amendments, SLAs including email trails between the client and the IT Service Provider. This latter point of email capture is increasingly used to highlight any communication that might indicate an implied contract variation by either party. I would suggest the inclusion of a Message Record Management (MRM) system as part of the KM solution.

Also, it will be necessary to install an activity monitor to log and track activity of users of the KMDB segment. In reality, this would be good practice across all of the KMDB segments but essential in this instance.

"However, be aware that a consistent under performance against SLA targets could be a fast track to IT outsourcing".

One final point. Where the service provider is internal to an organisation, for example the public sector, the risk of litigation is negligible. However, be aware that a consistent under performance against SLA targets could be a fast track to IT outsourcing.

I've seen this happen on a number of occasions. Although this is usually presented as an exercise in cost saving, invariably it is driven by a long term dissatisfaction in the performance of the internal service provider.

Cost of Compliance (Hot Spot 3)

Here is another example of the importance of a KM sub-set of material that can be assembled on the basis of a specific demand. During a compliance audit, ISO27001 for example, there will be a specific document set that will need to be made available to the auditors for the certification process.

Without a rigorous KM approach there is the risk of auditors finding a shortfall in the control objectives and controls. This will result in low auditor marking and possible non-compliance. There is now a real cost involved with the remedial work needed for a re-run of the audit, particularly with the high daily rates charged by external auditors.

The material can range from Information Security Policies to Physical and Environmental Security. There is a wide range of different types of information and data and the Audit and Map stages of the KM Framework will require a lot of research and agreement from the KM Stakeholders on what should be included in this KMDB Compliance segment.

It is likely that some of the lower level information may be located in Layer 2. If this is the case then it might make sense to leave it where it is and simply connect between the two layers. It is also true that the scope of ISO270013 is such that the KM will need to connect to a wider range of tools and assets.

One particular example is software asset management (ISO 27001 - Clause A8: Asset Management). Under this heading auditors will check the number and validity of software contracts held and check that the licences cover all the users who actually use the software. This could be addressed by setting up a source DB within a SAM tool and extracting all the data needed for the audit (as a controlled set) and then sending it to the KMDB. This is actually a very common failure point.

Risk to Service Continuity (Hot Spot 4)

In this final scenario I want to look at how the KMDB can be used to support Service Continuity. This has a much broader scope than just KM and I'm not intending to cover the whole subject of Business Continuity Management (BCM). Again, there are multiple terms involved here – like Disaster Recovery, Business Recovery and Service Recovery. In the case of ITSM and KM, I'm going to describe how KM can be used in support of Service Recovery within the broader BCM that covers the end-to-end business of a client.

The dilemma facing an ITSM organisation is no one can really identify all the situations likely to occur. Certainly, the evacuation of a data centre due to fire and flood is an obvious scenario, but thankfully not one that occurs very often. Clearly you can't prepare for every instance but it is possible to target some common 'knowns'.

So, here is a possible starting point. In our Layer 1 (High Business Value) under the Service Continuity category, the sub-categories should be constructed to reflect various 'threat scenarios' – one per sub- category, such as cyber threat, data theft and denial of service to name a few. We could also add major software outages that can and do occur from time to time.

Each 'threat scenario' can then be structured along the scope and guidelines of ISO223014. This will create a consistent framework for compiling all the recovery procedures, communication escalations and fall back plans for each scenario. Clearly, there is much more to discuss here but there is a future White Paper that will address all of these aspects of service recovery which is planned for publication later in 2015.

Conclusion

What this paper attempts to outline is a number of possible solutions to common issues around both risk and cost control in an ITSM organisation. It is not intended to be prescriptive. The KM system described here should be considered an 'entry level' system, but with the capability of extension as time and budget permit. This KM system is also predicated on content being held within existing repositories, as well as a central KMDB, but extracted on demand. The success of implementing a KM system will always reside with the management and staff of an ITSM organisation and not the technology. Hence the emphasis must always be on developing a KM Framework as the starting point.

The methods and techniques described in this paper are based on the KM service offerings available from CIH Solutions. For more information please contact Chris Hodder at info@cihs.co.uk. Also visit www.cihs.co.uk

Biographical Note: The author is currently a partner consultant with CIH Solutions and can be contacted at info@cihs.co.uk.

References:

1 ISO25964 Parts 1 and 2 – Information and Documentation - Thesauri and Interoperability with other vocabularies
2 Practice Direction 31B – Disclosure of Electronic Documents – Ministry of Justice.
3 ISO 27001 and ISO 27002 – International Standards for Information Security Management Systems.
4 ISO 22301 – Business Continuity Management Systems.