15th ICCRTS
"The Evolution of C2"

A Framework for Warfighter Information Services - using the concept of a Virtual Knowledge Base

Topics:
Topic 9: C2 Architectures and Technologies
Topic 3: Information Sharing and Collaboration Processes and Behaviours
Topic 2: Networks and Networking

Author:
Simon Bray, Dstl

Point of Contact:
Simon Bray
Defence Science and Technology Laboratory (Dstl)

Information Management Department
Rm C003, Grenville Building
Dstl Portsdown West
Fareham
Hants
PO17 6AD
United Kingdom

Tel:
+44 (0) 2392 537851

E-mail:
sebray@dstl.gov.uk

Dstl reference number: DSTL/CP40485

This page is blank intentionally (reverse side of cover sheet).

Title: A Framework for Warfighter Information Services - using the concept of a Virtual Knowledge Base.

Abstract: This paper presents a vision of how distributed information services and distributed knowledge stores can improve military effectiveness. Central to this vision is the concept of a Virtual Knowledge Base (VKB) – this includes machine reasoning capability, and supports human-machine teaming. The paper provides a Framework detailing the different kinds of information service that need to be provided and the different kinds of information that these services should act upon. These services range across Information Extraction, Information Collation and Validation, Situation/Intent Assessment, Situational Awareness, Decision Support, Collaboration and Process Support, and various supporting services which provide the technical basis from which to build and manage a Service Orientated VKB architecture. Implications, benefits and challenges of this approach are discussed.

1. Introduction

This paper is based on research undertaken by the author at the UK Defence Science and Technology Laboratory, which is developing the concept of a Virtual Knowledge Base (VKB) to support the Warfighter. A VKB will contain Information, "Known-facts" (defined below), and Warfighter Information Services – all distributed across a network. This paper proposes a suite of Warfighter Information Services which will process the information/known-facts in various ways and allow users to be made aware of them, explore them, and make decisions based on them. The key point is that this would, for the first time ever, make available to decision-makers all the relevant information/facts that MOD possesses in its networks. This would be done in a manner that is not overwhelming, presenting information to users based on the key-facts of a situation, allowing users to access all the supporting information on demand[1], and supporting their need to work collaboratively, within MOD and with external partners.

The proposed Warfighter Information Services are presented in a "framework", which categorises the kinds of services that will be needed. The framework is intended to be inspirational not prescriptive: it is intended to inspire others to think of innovative new services that will fit into the general areas identified in the framework. The services described here have been identified as "common-services". This means that these services can be used in various pre-arranged combinations to support all MOD's military business processes (i.e. those in the domains of Intelligence, Operations, Plans, Logistics, CIS, Personnel, etc.), at and between all levels of command. There will also be need of services that are "specialist" rather than "common". The framework does not define what these are but it does define how these services should "plug-in" to the framework, and hence have access to all the information/facts and common-services in a Virtual Knowledge Base. This ability for service to "plug-in" to the framework will be established through a set of VKB Design and Governance Principles, which have been created but are not described in this paper.

2. Related work.

This framework has its roots in many different places. A part of the purpose of the work has been to find a "unifying" framework that will support the needs of many different communities - encompassing all military communities who need to use information to do their job. The first starting point was the JDL model of Information Fusion [1]. This is popular in some communities but proved to be unpopular in the intelligence community partly because of its focus on a bottom-up assessment starting from "observed objects", whereas HUMINT is often received about intangible notions and about the ensemble-behaviour of collections of actors (e.g. morale and intent) without deriving this from observations of individuals/objects. Another problem was that the JDL model was interpreted as being fully automated where as the need for human

[1] Subject to security authorisation.

analysis/assessment or human-machine teaming was seen as being vital. Another problem is that it was seen as a fairly static "system" which processed a (fairly) steady stream of data from sensors, whereas what is also required is a capability to focus collection and analysis on different topics. These may or may not be fair criticisms of the JDL model but they have influenced this paper to use different terms and add additional services such that (hopefully) the intelligence and other communities can all be comfortable that their needs are met. The scope of this paper is also wider than JDL was ever intended to cover.

Another starting point was a range of initiatives by US DoD and standards bodies (notably W3C and OMG) related to Service Orientated Architectures, including the US DISA's Net-Centric Enterprise Services [2], US Navy's CANES and ForceNet, and their equivalents in the Air Force, Army and Marines. The SOA approach has spread to many NATO countries and is central to the NATO Network Enabled Capability (NNEC)[3]. There is a broad consensus among nations about the "core-services" needed to make SOAs work (e.g. messaging, discovery, authentication …), which is leading to NATO standardisation [4]. These core-services form a foundation of this Warfighter Information Services Framework, but the main focus of this paper is what is built on top of this "core". Here there is a disparity between nations' approaches (and even between the armed services in the US), and this paper proposes a new approach that the UK MOD may choose to adopt. This approach proposes a set of common-services for use by UK Army, Navy, Air Force, Marines and Joint Forces, and across all HQ functions (J1-J6, N1-N6, A1-A6, G1-G6).

Within the UK MOD the concept of a Virtual Knowledge Base was first coined by the Joint Doctrine and Concepts Centre (JDCC) (now called DCDC) in what became their "Inform Interim Concept" [5]. This spawned UK research on the VKB concept within the Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR) community. In parallel the Command, Control and Information Infrastructure community sponsored research on an "Information Layer". This paper attempts to unify these two previously separate strands of research and to recognise that the VKB concept is not only applicable to the ISTAR community but to all military functions.

Certain parts of this framework have been inspired by work in DSTO Australia on agent-support to Situational Awareness, exposed to the author through the US/US/CA/AS/NZ Technical Cooperation Program (TTCP) C3I Technical Panel 1 (Information Fusion) whose contribution is hereby gratefully acknowledged.

A companion paper [6] describes in greater detail the roots of the notion of "known-facts" and the use of machine-reasoning based on such "facts" in some but not all Warfighter Information Services. It also explains that this notion is similar to the W3C's semantic-web, and in particular RDF [6], but that the military requirements for expression of "known-facts" are richer than requirements for describing facts about resources available on the Internet (RDF's original purpose), and so a new formalism for the expression of facts is needed (and provided in [5]).

3. Virtual Knowledge Base Concept

A VKB is *Virtual* in the sense that the Information, 'Known-facts' and the Information Services in it are distributed across 'the network' in a way that is transparent to users, such that they do not need to know physically where information are stored or services are executed, and they can access the information and services from any point of attachment to the network (subject to security access rights).

A VKB is *a Knowledge Base* in the sense that it contains three kinds of knowledge:

- Knowledge of what is, was, and is expected/planned (here called *known-facts*[2]), stored in machine-understandable form (formally called "propositions") using an Entity-Attribute-Relation model. The scope of what is stored is specified through user-defined ontologies[3], which define the entity-classes, attributes and relations that are of interest to them.

- Records of all *information* that is input, created and used for reference purposes (and keywords/meta-data about these used for information retrieval/filtering etc.).

- "Know-how" made executable through *Warfighter Information Services* that act on the Information and Known-facts. Some of these services will employ machine-reasoning.

The relationship between facts, known-facts, information, and Warfighter Information Services is illustrated in Figure 1.

In addition the VKB contains *ontologies, policies* and *rules* that control the operation of the Information Services and define the semantics of the input/created/reference Information and known-'facts'. These are to be defined by MOD, and are to be adjustable and extensible at run-time (subject to appropriate authorisation controls).

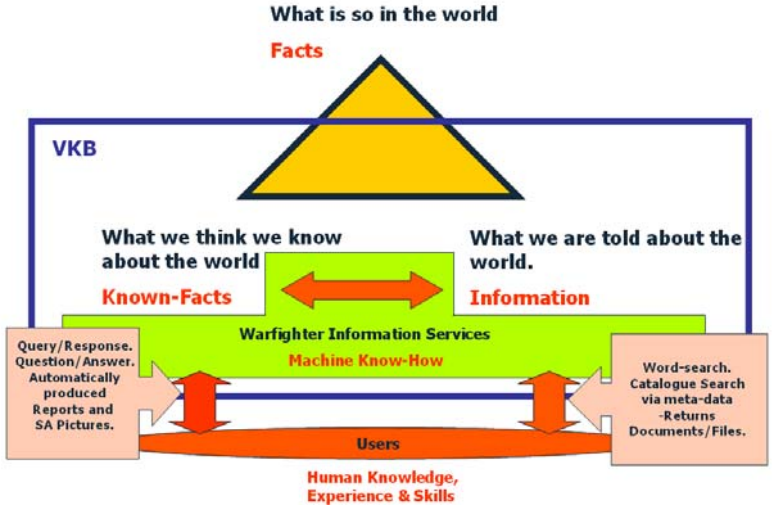

Figure 1 - VKB Capabilities

4. Warfighter Information Services Framework Overview

Figure 2 illustrates the framework as a whole. It comprises categories of Warfighter Information Service (the blue boxes around the outside), all acting on Information and Known-facts distributed across a network. It is important to note that this Framework is not organised along the traditional military function lines (e.g. Logistics, Intelligence, Command Inform and Battlespace Management, CIS/ISTAR Management, etc.). Rather the Framework seeks to establish a common set of Information Services that act on types of Information and Known-fact that are present in all military functions if abstracted in the manner described herein. The framework provides a set of *common informatic building blocks* from which different military functions can build the Business Processes that they need in a Service Orientated Architecture. The tools to build such Business Processes are part of the Framework.

[2] These so-called 'known-facts' may of course not be correct, i.e. the stored representation of facts may not correspond with reality. The term 'known-facts' has been coined to distinguish known-facts from facts, the latter being how the world really is.

[3] Note that these ontologies define concepts not literal terms that appear in text. These concepts are independent of language. Mapping from inputs in various languages to this abstract conceptual form is the task of the fact-extraction service described below.

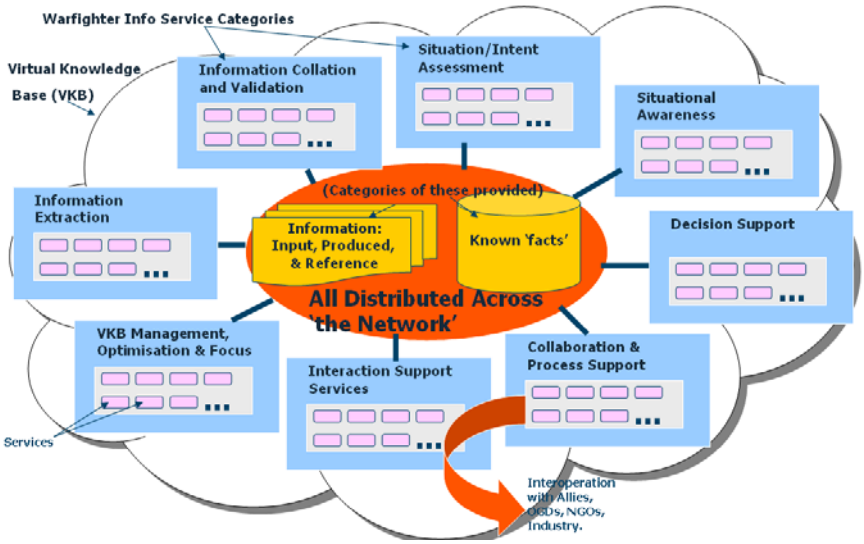

Figure 2 - Warfighter Information Services Framework

5. Information Extraction

Information extraction services are needed both to extract 'known-facts' from the input information and to catalogue and index the information to facilitate subsequent searching, filtering/browsing, analysis, and sharing.

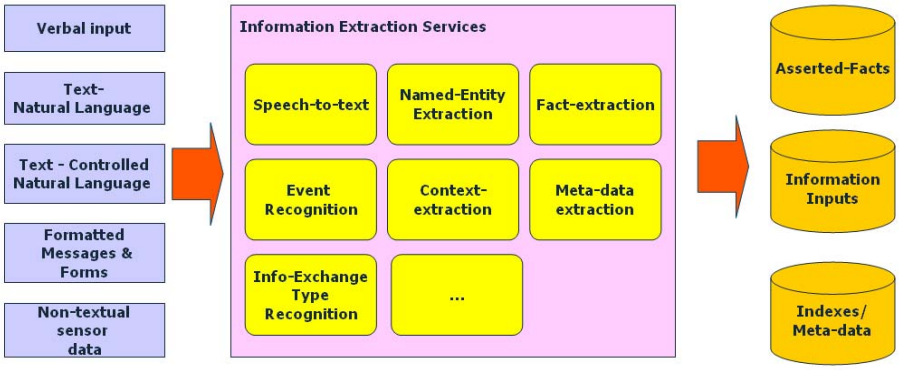

Figure 3 - Information Extraction Services

The *Named-Entity Extraction* service will process text in a natural-language or controlled natural language to identify nouns, or noun-phrases, that are instances of a set of entity-classes defined in a schema. The Named-Entity Extraction service being defined here does not simply identify named entities (or attach XML mark-up around them identifying the entity-class – which is what most Named-Entity Recognition tools do), it will also extract them into a separate index, with a cross-reference to the information-object from which they were extracted (and the position within it - character number of start and end). These then become meta-data through which it is possible to locate information-objects of interest by searching[4] for combinations of entity-class labels and their literal values.

[4] Search services appear in the Interaction Support services part of this Framework.

The *Fact Extraction from textual inputs* service will extract the *'*known-facts' that a community wishes to know[5] (N.B. not all facts) that are asserted or declared in textual inputs or inferred from them. Known-facts state that:

- An instance of a certain entity-class is said to have a given attribute-value predicated on it.
- An instance of a certain entity-class is said to have a given relation to another instance of the same or different (usually different) entity-class.
- Distinguishing whether the fact is concerning the present, past or future, and whether it has been asserted or declared.
- Plus various qualifications, attitudes (of someone toward a fact or set thereof), references, context, confidence, accuracy and other 'supporting information' about the cross-referenced *'known-facts'*.

It is important to realise that extracting *known-facts* is different from extracting *words*. What is being extracted here are the concepts that the words stand for, and the logical relations between these, which is independent of the language in which the words are expressed.

The *Event-Recognition* service will process streams of input-data and recognise that something of significance that is of interest to the community using this service has happened (i.e. an event), and to record this new 'known-fact'.

In the *Context-Extraction* service, we consider context to be a multi-faceted representation of the background to each information-object being processed. Unlike meta-data extraction (q.v.), context is not extracted from the information-object under scrutiny (the foreground). Context is partly discovered, partly sensed, and partly told. The context-extraction service maps every information-object that it knows about to a context-identifier, and stores facts about each separately identified context. This will then enable other services to (a) obtain a context-ID for a given information-object; and (b) query the VKB (not the context service) for facts about this instance of the context entity-class.

The *Meta-Data Extraction* service will analyse the content of an information object to determine various attributes of it.

The *Info-Exchange Type Recognition* service will analyse each Information-Object submitted for Fact-Extraction (and any associated Context-Extraction) to determine which type of Information-Exchange is associated with it. This service is needed to be able to specify business rules for handling each exchange-type.

6. Information Collation and Validation

Information and Collation can be performed manually, semi-automatically or automatically depending on the complexity/richness and variability/completeness of the inputs, the volume of inputs, and the criticality of the outputs to the organisation. The subsections describe the fully-automated case – as shown in Figure 4 - but it is up to each community to decide which parts it wishes to automate. This paper is not "recommending" full automation in every case.

[5] Which it specifies in an ontology.

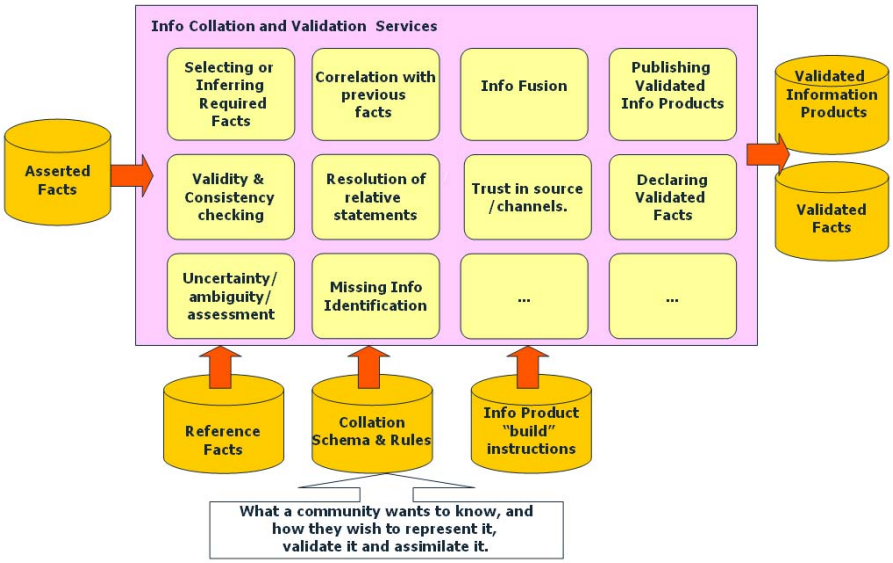

Figure 4 - Information Collation and Validation Services.

The Collation and Validation Services are taking as their input “asserted-known-facts” that are the result of the Information-Extraction services of the previous section.

The *Selecting or Inferring the Required Facts* service will take the newly asserted known-facts and try to extract from them what the community wants to know, which is recorded in its “Schema for Collation” (provided as an input to the service). How difficult this is depends on whether the “Schema for Collation” is a subset of the schema used to extract asserted-known-facts, or whether it has a different basis. If it is a subset then this service simply filters the asserted known-facts to select ones that match the schema for collation. If it is not a subset then schema-mediation services will need to be invoked to infer (where possible) what the community wants to know from what it has been told. In some cases this will not be sufficient, additional rules may need to be provided (as a stored-input to this service) to tell it how to determine the required facts from the inputs.

The *Correlation with Previous Facts* service will operate as follows. When new known-facts are extracted and asserted a new entity-instance-ID will be assigned to each entity that appears in the input information-object. This service will determine whether any of the previously known instances of this entity-class “match” (correlate) the new entity, i.e. it will determine whether the entity that is being reported is a new entity or one that was previously known about. Performing this correlation entails comparing some or all of the attributes and relations that have been predicated on the new entity-instance and previously known instances. Some attributes and relations will be defined as being essential for definitive correlation, whereas others will be indicative but not definitive. Similarly the values of some attributes will have an allowable “variation-margin” which is acceptable in a “match”. This variation-margin will arise partly due to observation errors, but also due to the fact that some attributes of entities vary over time (e.g. the location of moving platforms, but not of buildings). The result of this correlation may be definitive or probabilistic. It results in a new proposition that two entity-instance-IDs are equivalent (with any necessary qualifications about certainty).

The *Validity and Consistency Checking* service will perform various checks, including:

- That attribute values are of the type, range, and format declared in a schema.
- That the spelling of attribute values that are names (of people, places etc.) conform to reference facts (e.g. a Gazetteer) - the schema would not enumerate these names but would refer to the reference data to be used for validation.

- That an assertion of something being true has not previously been asserted as being false (or vice-versa). This is of course not necessarily an invalid situation – it may be that something was true and is now false. When these situations are detected the service must use rules (or heuristics) to decide whether the newly reported fact supersedes an existing fact (in which case the existing fact should be marked as 'historical' not 'latest and the new fact marked as 'latest'), or that there are conflicting of reports of what is so – in which case both facts need to be recorded as a mutually exclusive set of propositions.

In the *Resolution of Relative Statements* service, if a newly asserted fact contains a relative method of specifying its subject or object, for example, a location being: "100 m North of the xyz-named Hospital", then this service will resolve it into a definitive statement (a grid reference in this example). Other examples are, the subject being: "A's brother John"; or "the bridge over the xyz river near abc town"; or "the top of the xyz hill"; or "the day after tomorrow".

The *Information Fusion* service will operate as follows. As a result of the Correlation Service, it will be determined that some newly reported entity-instances are the same entities that have previously been referred to using different entity-instance-IDs (to some level of confidence). If this correlation has been established to be above a required confidence threshold then the Information Fusion Service will, where something new is predicated on each entity, assert new known-fact(s) using the 'old' entity-instance-ID but determining what is predicated on this entity-instance-ID from the 'new' known-fact(s) that have the equivalent/correlated entity-instance-ID. (Note the new facts are not destroyed by this process - they remain with their different entity-instance-IDs. This is so that the fusion decision may later be reversed if the correlation is called into question.) The result of this service is a set of new known-facts to be added to a VKB – the actual addition is done via the *Inform Service*.

Confidence/Trust Assessment is a heading under which come various services to assess the certainty and trust in the newly reported facts. These services will make use of the context-extraction service (q.v.) to retrieve a context-instance-ID, which they can then use to retrieve contextual facts pertinent to the new facts under consideration (e.g. how the information from which the facts were extracted were delivered, and whether the source was authenticated – which has Information Assurance implications). Other facts pertinent to this confidence/trust assessment may have been extracted from the foreground information, by the fact-extraction service (e.g. the originator of a message / document). The job of the confidence/trust assessment service(s) is to make a judgement on the basis of the available known-facts. If these services are to be common-services rather than bespoke specialist-services then a trust-model will need to be defined that is external to the service (as a set of business-rules that are expressed using labels that appear in a schema). This service then tries to apply the rules that appear in the trust-model. The result of this will be a declaration (or set of declarations) regarding the confidence/trust that has been determined through application of the business-rules.

The *Declaring Validated Facts* service will review the new facts asserted by the Information-Fusion service and make judgements about them using the outputs of the validity/consistency checking service and confidence/trust assessment services. If these checks and assessments are all positive then this service will declare that the new facts in question are valid. If there are reservations about the new facts they will be left as assertions and the results of the checking/assessment recorded as qualifications to the assertions.

The *Publishing Validated Info Products* service will take a set of "info product build instructions" as an external input and create a new information-product by following the instructions, which will entail transforming a defined set of validated-known-facts into a document, message or other form of information-object, which can be sent to humans or machines (distribution of info-products is a separate service).

7. Situation/Intent Assessment

Services under this heading enable human users to discover and judge what is going on in a situation, and to make assessments of intent, cause and capability. (These services can be applied in several domains including: Intelligence, Operations-Logistics plan coordination, CIS Fault Management, Cyber-defence …).

One approach is for humans review the available "information-objects" that have been collected and any information products created automatically from validated known-facts. This review will make use of the "Interaction Support Services" detailed in a later section. In parallel humans may use the 'known-fact-based' services defined here. These services perform machine reasoning of various kinds over the set of 'known-facts' in a knowledge-base, making use of Knowledge-Base Interaction services, and the results are then presented to humans through the User-Interaction Services.

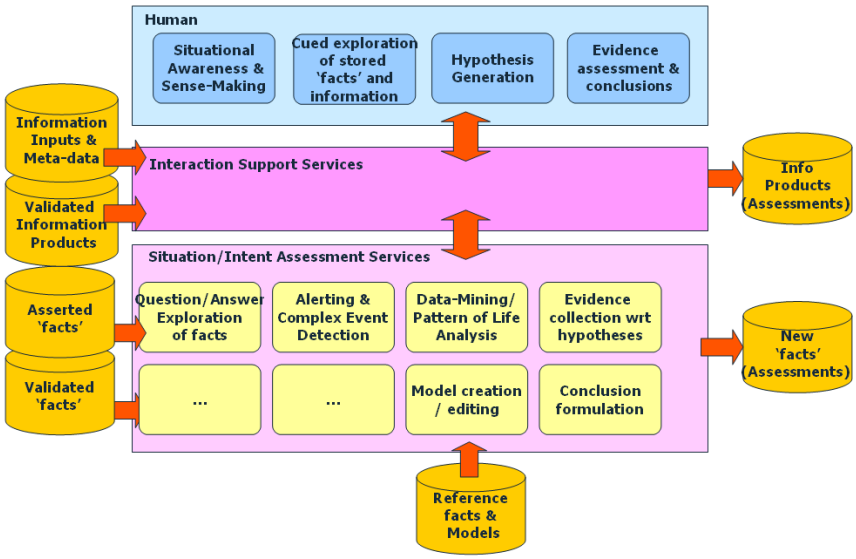

Figure 5 - Situation/Intent Assessment

The lower part of Figure 5 shows the services that have been identified to-date under the known-fact-base approach, and the upper part shows the human activities that these services are supporting. The intention is that the human and the machines providing the services act as a team, capitalising on the strengths of each party.

These services process the asserted and validated known-facts from the "information extraction" part of the WIS Framework, and through an interaction with the Human users who are responsible for the Assessment activity, generate new known-facts that represent the conclusions of the joint human-machine assessment.

The *Question-Answer exploration of facts* service will respond to questions posed in a controlled natural language, returning answers in the same form. This will entail use of the Controlled Natural Language Interaction service and the Question-Answering service. It will also entail remembering the sequence of questions and answers, such that the user may "drill-down" without having to repeat the previously established context for each new question (Such drilling down is signalled linguistically by words such as "the x", "this/these" etc. which refer back to what was previously discussed.) Only questions of the form "What, which, where, when, who, or how many" may be answered by this service – not "why or how".

The *Alerting of Complex Events* service will provide a user with an alert when a complex event that the user has previously specified occurs. (For example to be alerted when two individuals have been reported to be in the same location at the same time or within a specified time interval of each other – either in one report or in separate reports one on each individual.) This service is an orchestration of other services, namely: the *Complex Event Recognition* service; the Show service (to articulate an alert message announcing the occurrence of the event; publish-subscribe management and delivery services (to send the alert message to whoever requested it); and one of the User-Interaction services to present the alert in the required modality.

The *Data-mining and Pattern of Life Analysis* service will be used in conjunction with the *model creation and editing* service to help human analysts discover patterns of relationships and patterns of behaviour over time which are of interest. Being "of interest" will be defined with respect to a model which defines the classes of entity and their attributes and relations that are of interest, and which defines patterns that are of interest (either explicitly or by rules/heuristics by which patterns can be judged to determine their interest).

The *Evidence Collection with respect to Hypotheses* service will first assists a human user to form hypotheses in a Controlled Natural Language and then translate them into a set of hypothetical known-facts. Then the service will use the Distributed Inferencing Service to determine if any of the hypotheses can be inferred from any set of the known facts, either categorically or probabilistically, or that they can be refuted in either manner. If a positive result of either kind is returned then the relevant known-facts will be articulated to the user via the Show-service, and a User-Interaction service. If this exercise fails, this service may also report which facts are missing which would enable a positive result to be determined.

The *Conclusion Formulation* service will enable a human assessment to be first expressed in a Controlled Natural Language and then be converted into a set of 'known-facts', which will then be declared as being true on the authority of the human assessor (their truth may be qualified in various ways, including probabilistically). The service may, or may not, prompt the user as to what the conclusion should be, but as a minimum it must compare the proposed conclusion with the known facts and point out any contradictions to the human assessor. This service is an orchestration of several other services, as a minimum including: Controlled Natural Language Interaction; Fact-Extraction; Info-Exchange-Type recognition; Inform, Fact-Query, Distributed Inferencing, and Show Service.

8. Interaction Support Services

This set of services handles user-interactions, provides certain core-services necessary to support service interactions in an SOA (in a secure manner), and provides services to enable interaction Information-Objects and Knowledge Bases (containing *'known-facts'*).

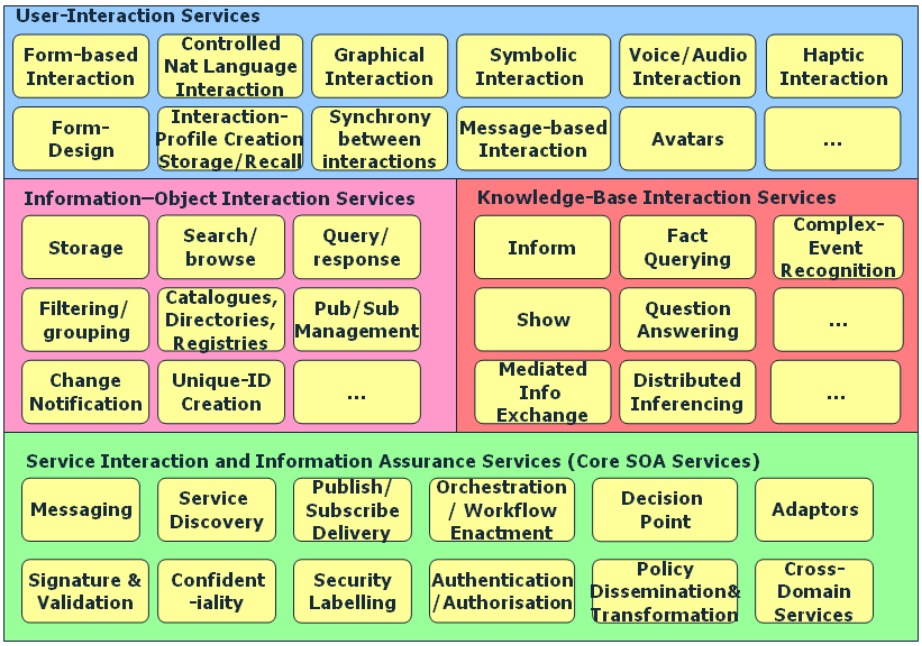

Figure 6 Interaction Support Services

User-Interaction Services. These services will allow users to interact with the Information-Object Interaction Services and Knowledge-Base Interaction Services shown in Figure 6, and interact with many other services which appear under all the other sections of this Framework (e.g. SA services). The services listed here enumerate the different modalities of user interaction, through: forms (as documents or web-pages), controlled natural languages, graphics, symbols, messages, movement of hands/body/head/eyes, speech/sound, and through Avatars). A 'profile' service is also required to enable Users/Groups configure how they wish to interact with the services in this framework, and to be able to store and recall different configurations to support different user-activities and different assemblies of computer-related hardware/software (for individual and group use; co-located and dispersed).

Which-ever modality is used (and they will normally be used in combination not singly), user 'interaction' entails: presenting information to users and enabling the users to navigate, 'drill-down', select information objects and invoke processing or presentation functions on them, and input information. A synchronisation service is required to ensure that when items are selected or changed using one interaction channel then these selections and changes are reflected in all the other interaction channels that are open (whether these have the same or a different modality).

Information-Object Interaction Services. These services enable users to have the following interactions with Information-Objects: storage (i.e. make persistent in the network), finding (through named-entity-search, or meta-data search, or via catalogues, directories and registries, querying (data-bases only), the setting-up of subscriptions to receive published information products, and grouping (for the purposes of filtering/browsing, or release control). These services will in part make use of the results of the meta-data extraction and named-entity extraction services in the Information Extraction section of this framework. The Catalogue, Directory and Registry services will provide facilities for making new entries, changing entries and deleting entries. All these services need to be capable of operating over a distributed network of information-object storage nodes, and to do this efficiently they need to operate in a *federated* manner, i.e. distributing their processing to coincide with the location of the distributed information-objects.

Other Information-Object Interaction services are provided to enable the above services and other services that use information to operate in a controlled manner. The Unique-ID service which generates globally unique alphanumeric identifiers, which are used by several services.

Knowledge-Base Interaction Services.

The *Inform Service* will assimilate new 'known-facts' into a knowledge-base as propositions, after performing the following checks/modifications:

- Distinguishing facts that are *asserted* and *declared*, and checking the authority, authenticity and integrity of *declarations*.
- If a new 'known-fact' is marked as being "the latest"[6] then the service will remark any previously recorded fact about the same entity as "historic".
- Check whether or not a new fact that is stated Categorically[7] contradicts previously known Categorical facts, and if it does then the set of contradictory propositions will be identified as a mutually exclusive set.
- Check to see if new facts resolve ambiguities that were previously noted (and recorded as mutually exclusive sets of known-facts). If the new known-fact is asserted as true that belongs to such a mutually exclusive set then known-fact previously recorded as possible will be set to "true" and the other possible alternatives set to "false". If the new known-fact is "false" then the action is simply to discount this option by setting the previously recorded known-fact to "false" instead of "possible".
- Apply any externally provided Business Rules which define how the new fact(s) affect what is already known, in the context of the current Information-Exchange session.

The *Show Service* will articulate selected 'known-facts' as textual messages, or as completed forms, HTML web-pages, geographic-overlays, symbols, text for avatars to speak etc. Inputs to the service specify which sets of facts are to be articulated, and the modality/language of presentation to be used. This is the opposite of fact-extraction: it is how what is known in a knowledge-base is communicated to users in a form that is easy for them to understand.

The *Query-Service* will determine if a set of postulated known-facts "matches" what is known, across the whole distributed knowledge-base. If the postulated known-facts contain variable names then the service will identify known-facts that match the input parameters that were not variable and then return the values of the variable parameters from these matching known-facts (or a fail response). If the input propositions did not contain any variable names then the response will be True or Fail (fail meaning 'not known'). This service needs to operate in an efficient manner across a distributed knowledge-base, and this will entail the service itself being distributed with the processing taking place at the locations that will minimise loading on the network to determine the results.

The *Question-Answering Service* will interrogate knowledge-bases and return answers to the questions. The person or machine forming the question need have no knowledge of how information is stored in the knowledge-base and he/he/it may phrase the question in a Natural Language or a Controlled Natural Language, and get responses in the same form. Answering a question is much more complex than responding to a query, though part of the process of question answering may be the automated formation of data-base queries and the synthesis of their responses into an answer to the question. For this to be possible at all there needs to be in place an appropriate ontology for the representation of knowledge relevant to the subject of the question, and this ontology must to some extent be shared by the question-poser and question-responder. This sharing of "mental models" may be tacit or explicit. (This is not the same as knowing what the "field names" are in any supporting data-bases.) Furthermore the response

[6] Every "known-fact" (a.k.a proposition) is marked as Historic, latest, future.
[7] Every "known-fact" (a.k.a proposition) is marked as Categorical, Probabilistic or Alternative.

mechanism must contain a reasoning capability that ranges over the ontology, enabling it to interpret the question in the context of the ontology, to map the information contained in supporting data-bases to the schema, and hence to determine which data needs to be retrieved and how it should then be processed to produce the answer.

The *Distributed Inferencing Service* will generate additional known-facts from facts recently added via the Inform Service by the application of a set of inference axioms (including in their domain everything relevant that was already known), and to do this in an efficient manner where the facts are distributed across a network. Conversely operating backwards, this service will determine if a postulated fact can be inferred from what is already known by application of a set of axioms (again in an efficient manner, distributed over a network).

The *Complex Event Recognition Service* will recognise the coincidence of multiple events – an event being the fulfilment of a condition defined over a set of entities; and coincidence meaning occurring at the same time, or in the same place, or time-and-place (time and place being extensions with specified bounds); or in a given sequence. Again efficient operation over a distributed knowledge-base is required, with recognition of the component events being performed local to where the pertinent facts are stored. (This service supports Alert Services and other services defined in other sections of this Framework.)

The *Mediated Information Exchange Service* will act as a go-between mediating exchanges of Requests, Commands, Proposals, Approvals/ Authorisations, Transactions[8], Offers/Bids, Polls (and perhaps other types of Information Exchange) between external parties, passing via a VKB, which "adds-value" in the process. Mediation may entail determining where to send requests or commands, and how (which communications channel, with what QoS/Information Assurance), and in what form/format (converting as necessary), keeping auditable records of messages sent and received (to support non-repudiation), etc.. It may entail brokering (matching offers to bids), and the checking of authenticity and integrity. Mediation may also entail keeping track of the state of patterns of message-exchange (e.g. request-made, awaiting response … accepted/counter-proposal made …), and associating each new message with the correct set of exchanges. In addition Business Rules may be provided to the service to enable it to make changes to known-facts in a VKB based on human decisions/declarations contained in the message-exchanges. Inputs to the Mediated Info Exchange service will be sets of known-facts extracted by the Fact-Extraction Service, which will also recognise and label the *type* of Information Exchange. Outputs will be sets of facts which are then articulated as an Information-Object via the Show Service, and delivered by the messaging service.

Core SOA Services. These services enable all the other services to interact, and to do so in a manner that provides the requisite Information Assurance. Details have been omitted as they are all well-known.

9. Collaboration and Business Process Support

See Figure 7. Descriptions of these services have been omitted to comply with the paper page-limit.

[8] A transaction here means a set of commitments entered into by external parties, which are intended to be represented as a set of changes to facts within a knowledge-base (these changes having external significance, including but not necessarily financial significance). This set of changes must be made atomically (i.e. all or none, and if a part fails all must be retracted).

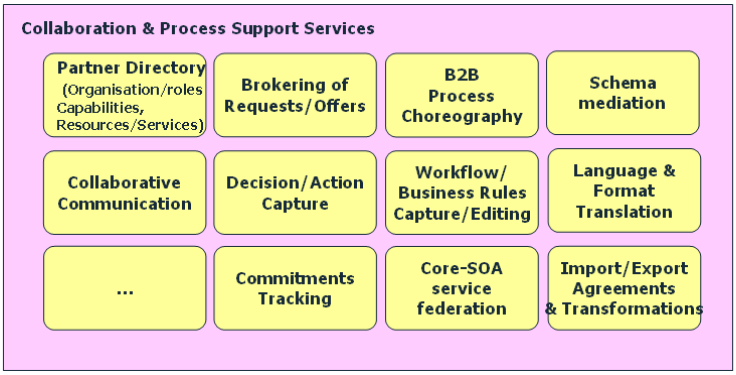

Figure 7 - Collaboration and Business Process Support Services

10. Situational Awareness Services

Situational Awareness is a cognitive state. Situational Awareness services provide information to Warfighters to help them to: make sense of situations in all its dimensions; be aware of threats or potential threats (abnormal behaviour); recognise the need to act and the opportunity to influence; be aware of the status of plans and resources; be aware of effects that are occurring and whether or not these were intended; and be aware of the overall status of an operation or mission.

The Warfighter requires Situational Awareness in many “domains”, i.e. covering different kinds of subject matter. These are shown on the right hand side of Figure 8.

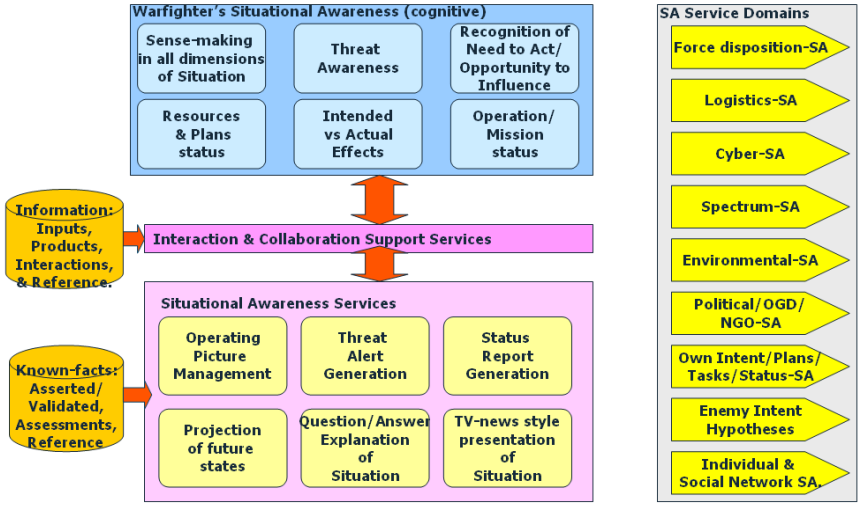

Figure 8 - Situational Awareness Services

Humans may enhance their Situational Awareness in many ways, not least by interacting with one-another, and reviewing and discussing information that is input or published as human readable “assessments” by an expert community. All these kinds of interaction are achieved using the Interaction Support Services, and the Collaboration Support services.

In addition Humans may draw on the specific Situational Awareness services described below.

The services in Figure 8 perform machine reasoning of various kinds over the set of ‘known-facts’ in a knowledge-base, making use of Knowledge-Base Interaction services, and the results are then presented to humans through the User-Interaction Services. The ‘known-facts’ that these services operate on are mix of asserted facts extracted from externally supplied

information-objects, validated facts, assessments declared as facts, and reference facts (e.g. locations of place-names).

The *Operating Picture Management* service creates a 'view' of the currently known-facts and presents this to users. The specification of which known-facts are to be included, and how they are to be articulated will be defined as an information-object which persists independently of the service but can be edited via the service. Many different 'views' can be defined and called up by different user-groups or individuals as required. The service will then instantiate selected views at run-time, updating the views in near-real time as relevant facts change, calling upon a range of other services to achieve the desired effect.

The *Threat Alert Generation* service will alert subscribers to this service that there is threat to them, based on their currently known location (in the VKB), and a user-specified radius of concern (which could be threat-dependent). The details of how the service is orchestrated from other services is similar to the Operating Picture Management service, i.e. there would be an editable specification of what threats are of interest, these would then be queried, filtered by location of each subscriber (determined from the subscriber's identity via a knowledge base query), be articulated in a specified manner, and then be delivered and presented to subscribers to the service in their own personalised modality (which could well be voice in this case).

The *Status Report Generation* service will provides on-demand, event triggered, or periodic reports of the status of specified entities, and their relations to other entities. Status reports can be simply about the world as it is reported to be; or can be comparative against a plan. When comparing reality with plans (either plans that are options or plans that are committed to execution), the assumptions in the plans can be compared with the known-facts (which will often entail use of the *distributed inferencing* service). For plans that have been committed to execution the current disposition of forces and resources can be compared with the planned disposition and exceptions reported. Status reports highlighting variances can be provided in any required modality.

All three of the above services are similar in their internal composition, what distinguishes them is: (a) under what conditions they deliver results; (b) the nature of the facts being reported – which is immaterial to the design of the service.

The *TV-news style presentation of a situation* service will create a "news-story" on demand or triggered by events, and present it using a combination of spoken text (with or without an avatar), written text, geographic overlays, images/video, etc. all synchronised to illustrate the spoken story[9]. This goes beyond presenting the "current situation" which the previously mentioned services do. It entails determining a sequence of "significant events" over a period of time, presenting evidence that these events have occurred, and articulating an assessment of how the events are linked and what this pattern of behaviour may indicate. The assessment is not made by this service - it uses the assessments stored as 'known-facts' that are created by the Situation/Intent Assessment services (which will often involve human-machine teaming).

The *Question-Answer Explanation of a Situation* service will respond to questions posed in a controlled natural language, returning answers in the same form. This will entail use of the Controlled Natural Language Interaction service and the Question-Answering service. It will also entail remembering the sequence of questions and answers, such that the user may "drill-down" without having to repeat the previously established context for each new question (Such drilling down is signalled linguistically by words such as "the x", "this/these" etc. which refer back to what was previously discussed.)

[9] This service has been demonstrated in a research environment by DSTO Australia, and information about it provided to Dstl via The Technical Cooperation Program C3I Group, Technical Panel 1 (Information Fusion).

11. Decision Support Services

The Decisions being considered here are made by humans. Decision Support services provide information to Warfighters to help them to: formulate courses of action in accordance with higher command's intent, generate plans, decide on a course of action and issue orders to execute a plan, and perform dynamic re-planning during execution.

Decision Support is required in many "domains", including Command and Battlespace Management (CBM) in different operating environments, and provides Command and Control over different kinds of asset. These domains shown on the right hand side of Figure 8.

Humans may make decisions in many ways, not least instinctively and by interacting with one-another, and reviewing and discussing information that is input or published as human readable "assessments" or "plans" by an expert community. All these kinds of interaction are achieved using the "Interaction Support Services" and the Collaboration Support services.

In addition Humans may draw on the specific Decision Support services described below. These perform machine reasoning of various kinds over the set of 'known-facts' in a knowledge-base, making use of Knowledge-Base Interaction services, and the results are then presented to humans through the User-Interaction Services.

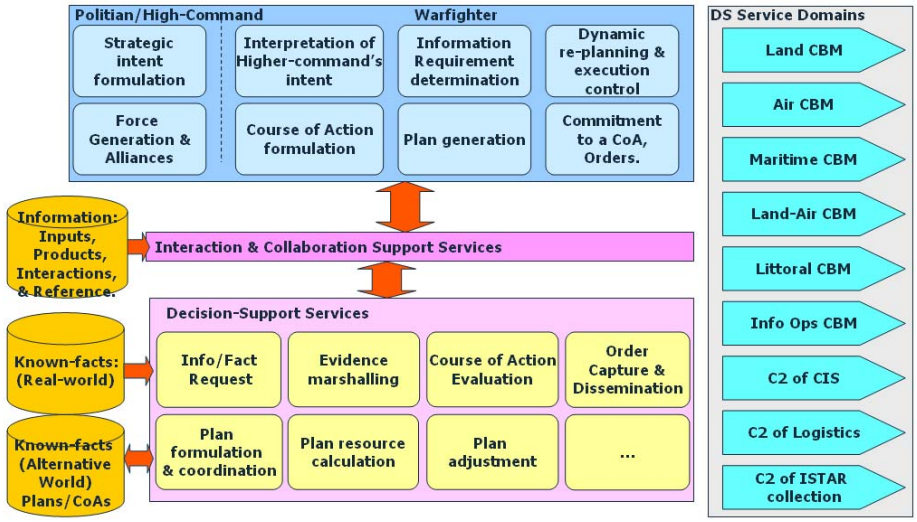

Figure 9 - Decision Support Services

The 'known-facts' that these services operate on are mix of: real-world facts, including asserted facts extracted from externally supplied information-objects, validated facts, assessments declared as facts and reference facts (e.g. locations of place-names); and "alternative-world-facts" i.e. plans and courses of action formulated as sets of known-facts, having in them representations of real-world entities, but are distinguished from real-world facts in a VKB to avoid confusing plans with reality.

The Decision-Support Services identified to-date are shown in the bottom part of Figure 9 and are described below.

The *Plan formulation & Coordination* service will enable users to capture their plans in machine-understandable form[10], and (if required) to do this on a large scale, with multiple planning teams (co-located or distributed over a network), and achieve coordination in multiple dimensions:

[10] The proof-of-principle that this can be done has been demonstrated in the UK-US International Technology Alliance on Network and Information Sciences. This project has developed a Coalition Planning Model [7]– which is allows the pertinent aspects of a plan to be represented in machine understandable form.

between different levels of a command hierarchy, between forces operating in a given operating environment, and between functional elements (Ops, Logs, Intell, CIS ..). Plan coordination through this service will include: checking of consistency of starting assumptions between parties and between each party's assumptions about what another party will put in their part of a plan and what they actually have put in their part of a plan (e.g. in relation to resources employed, intended effects, timing/routes of movements, timing/location of actions, etc.); brokering to achieve the most effective use of resources that are shared between parties; brokering of the scheduling and location of activities that need to be synchronised – either to achieve a joint effect or to reduce the risk of own-force causalities.

Plan coordination also comes about by making explicit the rationale for tasks, and sharing this understanding between the parties. Rationale will be captured as part of this service and each party's rationale made available to human operators as they review and develop plans.
All these coordination activities can be performed continually as a plan develops, not just at the end or major review point. This is one of the major benefits of a networked approach – plan divergence can be spotted early and corrective action taken in a timely manner.

The *Plan Resource Calculation* service will perform computations of resource requirements based on plan segments. This will require a model to be provided to the service which defines how resources are derived from plan entities, attributes and relations by a process of logical inference. The service will then allow users to select a portion of a plan for analysis (i.e. selecting a set of plan entities and a time interval), apply the model to these entities and report the result. Whilst the model is external to the service, the service will also provide the capability to review/ develop the model. The internal form of the model will be a set of executable logical statements which apply to the schema of knowledge representation (i.e. the ontology) being used for the plan. These logical statements can be axioms (stating what resources are always entailed) or can be conditional statements.

The *Plan adjustment* service is intended to support the need for dynamic re-planning after a plan is committed to execution. Users will make use of the automated Situational Awareness services to understand current status and threats, and how these relate to plans. By this or other means users will identify that parts of a plan need to be changed. This service will enable users to select a plan segment and to change assumptions and intended dispositions to match reported reality (by controlled importing of known-facts not by manual entry). It will then support "what-if" decision making in the following manner. The user(s) may wish to be advised:

- whether the original end-state can be achieved, to the original schedule, by an increase/decrease or re-distribution of resources;
- whether specified changes to the schedule are achievable within available resources and other constraints, and what the risks would be;
- what knock-on effect there would be if part of a plan were changed in some way specified by the user – and which of the original objectives could be met or would be put in jeopardy by such changes?

In such situations time is of the essence, and it will be vital to ensure that all affected parties are made aware the impact that plan changes would have on their activities, and either to seek their agreement to changes or command them to change depending on the nature of the relationship between the parties. This service provides the means to:

- Identify which parties would be affected by changes, and what the impact on them would be;
- Create one or more information-objects that describe the change, the rationale for the change, and the impact of the change in human readable form (which can be edited manually prior to release);
- Assist the user to send this information to the affected parties, using whatever type of Information Exchange is required (i.e. as a notification/warning, request for comment, request for approval/ order).

As and when plan changes are approved (by whatever means not necessarily using any of these services), all affected parties will use this Plan-adjustment service to commit the changes into their own plan segments (if authorised to do so).

The *Order Capture and Dissemination* service will assist users to draft Orders following any specified template (e.g. STANAG 2014: Warning Order, Operations Order, Frag Order, and their various annexes for Logistics, Comms, Movements, Aviation etc.) Where possible, and under user control/revision, the contents of the various sections will be partly filled-in automatically using alternative-world-facts from plans that have been developed using the plan formulation or adjustment services. When an order is complete the service will then instigate its dissemination to those that it affects, and keep track of acknowledgements received in response. It will not normally be necessary to send the whole order (with all its annexes) to all the recipients, and the service should facilitate sensible partitioning of the information for different recipients to minimise network loading (and information overload of the recipients). All recipients should however have the right, and the means provided by this service, to request and obtain other parts of the order in a convenient manner

The *Information/Fact Request* service will assist users to fill gaps in their knowledge of a situation. A user will use this service when he/she believes that more information/facts are needed to make a decision than are available. A user may request Information or Facts, or both. A request for a fact is made by asking a question, which is answered by facts or yes/no. For example: “are there any enemy forces in this area, and if so where?”, or, “is bridge at coord xyx still standing?” Whereas a request for information will return an information-object from which the recipient will need to make his own assessment (e.g. an image).

The *Evidence Marshalling* service will assist users to monitor evidence relevant to their decisions. The service will enable individual users and groups of distributed users to:

- Define what facts they need to know to make a decision in a structured manner (based on an entity-attribute-relation schema);
- Establish persistent queries which periodically refresh local knowledge of the identified facts;
- Set-up Info/Fact Requests (through the Info/Fact request service);
- Present what is known about the identified facts through user-defined modality of interaction. On-demand users should be able to drill-down to the information-objects from which the known-facts were extracted (if authorised and if bandwidth is available), and display all the known-facts predicated on these facts e.g. asserted probability, validity conditions, trust parameters, who asserted or declared it, etc.

The *Course of Action Risk Evaluation* service will capture alternative courses of action (at a fairly high level) and then enable users to evaluate the risks of different courses of action based on currently known-facts. This will require a model of risk to be provided to the service which defines how risks arise from entities in threat assessments/operational pictures in conjunction with entities in courses of action. The service will then allow users apply the model to the entities in the courses of action and report the results. Whilst the model is external to the service, the service will also provide the capability to review/ develop the model. The internal form of the model will be a set of executable logical statements which apply to the schema of knowledge representation (ontology) being used for threats, operational pictures, and courses of action. For example the model may contain a statement to the effect that if a friendly-unit is within a range of 5 km of a given enemy weapon system then it is at risk, or within 200m of an IED, etc. The service will then apply the rules in the model to the proposed course of action, with the known locations of enemy forces and weapon systems to establish a risk score, and a confidence level in the risk score. It may well be that certain information is not available, in which case a set of standard assumptions may be applied (which users can edit) – for example that certain enemy units have certain weapon systems.

12. VKB Management, Optimisation and Focus

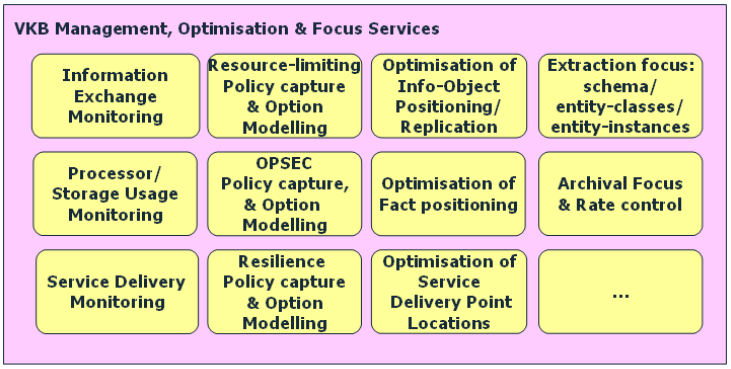

Figure 10 - VKB Management, Optimisation and Focus

These services ensure that:

- The operation of a VKB is observable and controllable.
- The total system can be optimised adaptively meet a number of requirements including minimising network loading and providing the requisite amount of resilience.
- It is possible to focus what knowledge a VKB contains in-line with operational needs.

The services in the first column of Figure 10 will provide basic monitoring services, which would be orchestrated with user-interaction services to provide human visibility of the various aspects of the VKB being monitored.

The services in the second column of Figure 10 will provide the levers through which control will be exercised over a VKB to ensure that it: (a) does not exceed resource limitations, including various means of allocating resources to communities, operations/missions etc.; (b) provides the requisite amount of resilience (through replication of information-objects and having multiple service-access points); and (c) provides the security appropriate to the current operational needs – including network security, info access control and info release control between communities/ collaboration partners (enforced by the *Cross-Domain Service*). In all three cases, a-c, this control will be exercised by defining policies. These services will enable military users to define policies at a high-level that is independent of the many technologies used to implement a VKB. These services will not only provide the physical means to capture policies in a simple fashion but will also enable the policy-setter to visualise the impact of policy-options, as each option will entail one or more trade-offs within a VKB. Thus these service will also include models that allow these trade-offs to be explored and an appropriate “operating point” selected by a user.

The services in the third column of Figure 10 will enable optimisation of the operation of a VKB over a network, i.e. optimisation of the way that information-objects, known-facts and services are distributed over a network. The optimisation will seek to optimise the placement of information-objects, known-facts, and service delivery points such that network loading is minimised within acceptable constraints on performance, and within the constraints specified through the policies from the services in the second column. The optimisation services will make use of the monitoring services from the first column of Figure 10 to learn how the VKB has been deployed across a network, how it is being used, and what problems are manifest.

The services in the forth column of Figure 10 will enable operational staff to control the volume and focus of knowledge extraction, storage and archival. This enables use of processing and storage resources to be focussed on obtaining, and having to-hand for immediate use,

knowledge about facts that are pertinent to current operations. This is both to ensure that a VKB does contain what is most needed, and to ensure that it does not become "cluttered up" with knowledge about facts that are of little or no interest thus wasting valuable resources.

The *Extraction Focus* service will enable operational staff to specify which schema are to be used for fact-extraction, and if desired to further constrain extraction to certain entity-classes within schema, or about individual instances of an entity-class.

The *Archival Focus and Archival Rate-Control* service will enable operational staff to specify how long *know-facts* should be held in a VKB that are marked[11] as either:

- "historical" (i.e. not "latest"), or;
- "predicted" but after the time to which the prediction refers.

This holding interval could be uniform or be varied in order to retain some *known-facts* longer than others, hence varying the time-focus of which *known-facts* are immediately available for access via the services of this Framework.

13. Implications of this approach

(1) Cost-reduction should be possible in the delivery of IM/IX capability, and capabilities that depend on this, namely Situational Awareness, Logistics, and Decision Support. This follows from:

- Recognising that there are "common-services" that are applicable across multiple capability areas.
- Reducing the "integration cost" through use of a common conceptual framework, and implemented within an Open Architecture that supports this conceptual framework.
- Reducing the cost of upgrade/refresh by allowing open competition on an incremental basis, made possible by a combination of an MOD-controlled Open Architecture and a service-based approach to architectures and acquisition.

(2) A vast improvement on current capability to access information and use it to military effect in a timely manner. (This follows from the VKB concept of information being "in the network", to which all users/platforms are connected could have access subject to their need-to-know and other security requirements.)

(3) A revolutionary new capability to exploit information, by harnessing the power of machine reasoning (in services) and Human-Machine Teaming to extract maximum military benefit from information and intelligence. This will result in:

- Better coordination of planning across J1-J6 cells and their single service equivalents to synchronise their effects, increase the tempo of their planning – and enable dynamic re-planning.
- Better awareness of threats (including IED threats) at the planning stage.
- Faster, safer and more-effective reaction to events. Including the automated dissemination of threat-warnings to whom-so-ever is discovered, from the available facts, to be under threat or at risk from enemy and own-force activities.
- Better understanding of complex situations - hence better decisions.
- Reduced human "information overload" (or smaller HQ staffs).

(4) Increased capability to manage information flows over networks such that:

- Flows are arranged in an efficient manner (for example not all going to a central repository and being replicated).
- An appropriate level of resilience is provided (neither too much nor too little, depending on the operational need).

[11] The marking historical/latest/predicted is a flag attached to every 1st-order proposition.

- Where there are insufficient network resources to meet all information flow needs then there are mechanisms provided that are sensitive to the operational priorities. This will also help achieve better military outcomes when networks are under attack.

(5) An increased ability to configure the available infrastructure (networks and information services and security mechanisms) in a "plug-and-play" manner, and hence delivering operational agility for the deployment of CIS.

14. Research Challenges

Some of the key issues that need research to realise this vision are:

- Consensus on a common-language for the expression of known-facts in communication between services of this framework.
- Fact-extraction from Controlled Natural Languages: robustness, ease of use …
- Question-Answering: reliability, speed, mental-model capture …
- Articulation and portrayal of known-facts to users: how to do selection, arrangement into a logical "story", portrayal in different modalities …
- How to operate services acting on distributed known-facts efficiently and resiliently over military networks in a near real-time fashion.
- How to optimise the distribution of services and information/known-facts to minimise network loading yet provide the specified level of resilience.
- Understand the impact of the VKB concept on network capacity requirements.
- How to manage the development and deployment of services and schema by multiple parties on a large scale.
- How to maximise the effectiveness of Human-Machine Teams.

15. References

[1] Joint Directors of Laboratories (JDL) model of Information Fusion, of which there are many variations, for example: A. N. Steinberg, C. L. Bowman and F. E. White, "Revisions to the JDL Data Fusion Model", in Sensor Fusion: Architectures, Algorithms and Applications, Proceedings of the SPIE, Vol. 3719, 1999.

[2] US DISA Net-Centric Enterprise Services (NCES). http://www.disa.mil/nces/dev_tech_resources.html

[3] NATO NEC Feasibility Study, NC3A, CIS Division, 2005.

[4] (NATO) Core Enterprise Services Framework (v1.2), Sub-Committee 5, Core Enterprise Services Working Group. 2009.

[5] The Inform Interim Concept – a paper by DCDC, 29th Nov 2007 (UK MOD).

[6] Human and Machine interaction with Knowledge Bases, Dstl/CP40411, S E Bray, 15th ICCRTS, Jun 2010.

[7] Coalition Planning Model, US-UK International Technology Alliance on Network and Information Sciences. D Mott (IBM) et al. See http://www.usukita.org – Project 12