Chapter 3

GLOBAL WORKFLOW STRATEGY

Creating and Managing Multilingual Knowledge e-Content

Nazli Choucri, Farnaz Haghseta, and Toufic Mezher

Introduction

In Chapter 2, we reviewed the design and implementation of GSSD, a knowledge e-networking system focusing on transitions toward sustainable development. Central to this effort are defining the knowledge-objectives, delineating the logic for knowledge e-networking, identifying critical e-barriers and designing solution strategies to help reduce structural impediments to knowledge access.

Operational Aspects of e-Collaboration

In this chapter we extend the collaborative strategy from knowledge e-networking activities, *per se* and examine the fundamentals of creating and managing a distributed global system. In so doing, we address conceptual, computational, and other implementation aspects of distributed e-networking across languages, regions, and cultures.

Given that our major objectives are to engage in and expand diverse forms of international collaboration around knowledge management and networking, GSSD seeks to help improve our understanding of the challenges and opportunities for sustainability. While the application pertains to sustainability, the connectivity issues are generic in nature, as are the complexities of international collaboration and the management of distributed knowledge systems.

The overarching challenge consists of the creation, maintenance, and management of a multilingual knowledge-base. In other words, the challenge involves defining the overall strategy, the technical requirements, and the management and organizational aspects of the e-networking process. The challenge is also evaluating a work-in-progress as we learn from earlier

N. Choucri et al. (eds.), Mapping Sustainability: Knowledge e-Networking and the Value Chain, 49–66.

experience, and seek to adapt, adjust, and otherwise enhance overall e-performance and collaboration.

More specifically, we highlight four anchors for global e-collaboration. First is the *provision of knowledge* in terms of content-submission to the knowledge-base. Second is the nature of the *workflow process*, namely how new knowledge content that is submitted in diverse locations worldwide is processed throughout the system – from the initial input, to web published status, in all language-databases and all servers in all locations. Third is the design and implementation of a *multilingual strategy* for distributed networking systems worldwide. And fourth, is examples of *customized* design responses are needed to address different situations on-the-ground.

The Workflow Process

The managerial and computational innovations reported in this chapter revolve around the creation of the workflow. The *workflow* is defined as the process through which new knowledge content is routed through the system. Specifically, the workflow process refers to the path through which each individual knowledge item is guided – from site selection, to submission of abstracts, to translation into GSSD's supported languages, and finally to publication on all GSSD mirror sites. The workflow consists of the process as well as the protocol governing the activities of all partners as they organize the knowledge content and its flow throughout the e-network.

To be effective, when a workflow involves interactions among members of a distributed e-network, it must be designed to enable customization as needed. One size seldom fits all. Accordingly, in this chapter we show how GSSD operates as a globally distributed endeavor by defining the diverse functions of knowledge provision, distribution, and translation, and then by specifying the steps that must be taken in order to ensure overall performance. In so doing, we provide the backstage view of what actually happens computationally when knowledge is transmitted from *Submit Site* and then routed from the initial provider to the GSSD system as a whole.

The issues addressed here are interconnected. On the one hand, we show the management and design issues central to the workflow process; and on the other we provide the basis for system-transparency – in case a reader wishes to know *how* the submission process is done, and *why* it is done this way. Thus, we provide the basic knowledge required for others to understand, assess, replicate and most certainly improve upon the actual *practice* of e-collaboration that we have developed.

Given that GSSD search functions operate across languages, and that the abstracts are rendered *multilingual* as well, the knowledge management process requires a certain degree of precision in rendering and coordinating

knowledge items across language databases. Since all languages are available in all mirror sites, at all times, the challenge is to maintain cross-location consistency such that the mirror properties are retained throughout the network.

In short, our purpose in this chapter is to render matters of *system design* and *operational strategy* for a distributed global e-system as transparent as possible. The challenges associated with such an effort are daunting, and a certain degree of trial and error – with much error – is a pervasive feature of the entire enterprise. Then, too, given that cyberspace is an English-speaking venue for a world where most people do not speak English, constructing multilingual capabilities is an especially relevant endeavor.

3.1 Selection and Content Provision

The value of knowledge lies in its content. In the absence of a robust method for valuing content, it is incumbant upon us to articulate the 'quality rules' for knowledge provision. This section begins with a discussion of methods and mechanisms for content selection, review, and submission into the knowledge-base. Most of the issues addressed can be seen as methodological and proceedural in nature, but they are at the very heart of any distributed knowledge system.

In this connection, recall that *Mapping Sustainability* in Chapter 1 provides the framing system for sustainable development as an area of knowledge. The knoweldge content therefore consists of systematic observations of the domains of human activities as well as the dimensions associated with each domain. Recall also that Appendix A, at the end of this book, presents at the 'Guide to Core Concepts,' in the familiar form of tables of content. This knowledge display serves as the bridge between the theory underlying this initiative analysis on the *conceptual* plane and the methdological features at the *implementation* level.

The issues addressed here pertain to the profiling of knowledge, to its abstracting and recording, and to the practical application of conceptual guidelines. Increasingly the international community as a whole, and the scientific, educational, and knowledge communities more specifically, are recognizing the imperatives of decentralization and the invariable attendant 'push' toward localization. More fundamentally, we show some of the practical, strategic, and essential tasks that accompany any collaborative knowledge-based initiative – particularly those of global scale and scope.

3.1.1 Identifying Content

To begin with, Internet-resources must be identified as relevant to GSSD's knowledge interests and overall goals, then specific features of knowledge are summarized into individual abstracts, and then these are classified according to the GSSD knowledge system index-structure and Guide to Core Concepts, and finally they are reviewed for consistency and quality prior to web publication.[1] In general, individual sites whose documents are abstracted for the knowledge-base must be available at no cost to the user, and, preferably, with no advertisements.[2]

Given that the focus is always on content, large sites need to be disaggregated into smaller, more focused content and then differentiated into further elements, as relevant.[3] Pages of links are useful, but they must be reviewed and can often be spidered. By following all the links to outside sites, it should be possible, except for the very largest sites, to link directly to the original sites rather than through a links page.[4]

3.1.2 Sources of Knowledge

The knowledge-base of GSSD consists of organization-supported *holdings* selected from the following types of sources:

(a) *Formal Institutions*, which include well known leading organizations focusing on science and technology, or institutions of governments and governance, business and industry, as well as major representatives for civil society (notably non-governmental organizations). We assume that the institution itself is responsible for assuring quality and reliability. In essence, we consider reputation and the existence of some accountability process at the source.

(b) *Content-Specific and/or Scientific Institutions*, which include the organizations described above, plus civil society affiliations, such as the Population Council, the Energy Defense Council, United Nations Development Program, Government of Brazil Ministry of Energy, etc.

(c) *Intermediation Service Entities* such as public service or private sector third party associations of institutions or publications, such as the

[1] The index structure refers to the details in the contents of the Guide to Core Concepts.

[2] Exceptions can be made if absolutely essential.

[3] In the case of a large site that is being disaggregated into individual topic-board sections, an abstract should still be created for the main page. (In the case of site reorganization, the links to smaller sections of the site may become dead links, but the address of the main page will most likely stay the same.)

[4] Often it is also useful to make separate abstracts for relevant sites. This practice minimizes the work required by the GSSD user to find the information he/she is looking for.

e-journals or e-reports of the World Business Council for Sustainable Development.

(d) *Specialized Submissions*, or e-materials created by individual researchers, scholars, and other non-profit professionals.

3.1.3 Coverage

The types of e-materials (i.e. knowledge content) to be considered for inclusion in GSSD include the following:

- Agreements
- Bibliographies, reports, and journals
- Case studies
- Definitions/theories
- Events
- Indicators/data
- Models
- Organizations

The GSSD Glossary of terms, presented as Appendix B at the end of this book, includes definitions of the above, as well as definitions of all key terms relevant to the GSSD knowledge content and its organization.[5]

3.1.4 Knowledge Management

Management of the GSSD knowledge-base consists of two distinct but related tasks. First is the organization and classification of knowledge items. Second is the periodic review and updating of the overall knowledge-base. This task includes the identification of *dead links*, namely, e-addresses no longer in operation, either because the address has changed, or because the content is no longer relevant. Once the dead links are identified, the GSSD team locates and replaces the materials as appropriate.

3.1.5 Organization and Classification

Two procedural and system features are critical operational guides for these aspects of content provision. First is use of the sustainable development

[5] Each candidate knowledge entry (abstract item) must have at least one of the following criteria: be distinctive and identifiable, provide some specific rather than general statements, except of course when dealing with policy statements; have some possibility for validation or replication; and point to specific content pages, except for abstracts linking to home pages. Excluded from consideration are daily news items, popular magazines, opinion papers, etc.

ontology derived from the framing system and the ontology rules – in terms of slices, rings, cells, concepts, sub-concepts, and the fifth ring, i.e. international response and global accords – to provide the semantic and theoretical frame for 'tagging' each abstract. Second is the methodological requirement that each *knowledge abstract* must, to the extent possible, use the content-driven differentiation at the intersection of slices and rings. This refers to the detailed specifications within each of the categories in the Guide for Core Concepts. Finally, the dominant bias in this process is that of *exclusion* – on quality and relevance grounds – with the corollary that inclusion is not an operational directive.[6] The strategy is to purposively maintain a streamlined knowledge-base.

3.1.6 Reviewing and Updating

By definition, a dynamic knowledge-base is always changing, and we would expect it to evolve over time. This is the case with GSSD as new materials are included and outdated materials are either updated or removed entirely. Since the knowledge-base consists of indexed and abstracted materials of e-resources already in the public domain, retaining clean files is an important challenge. At this point we turn to quality control guidelines – in terms of *what* to do and, equally important, what *not* to do – in order to avoid inadvertent problems in the process of adding content to the knowledge-base. These same principles also govern our decisions when making changes to existing content in the context of the discussion in Chapter 1. At issue here is the process through which the 'slots' in the framing system are 'filled' and the ways in which the extensions are made from individual knowledge items–to enhance the design and contents of a knowledge-base. Thus, in order to maintain, update, and revise the GSSD knowledge-base, and to assure quality and reliability in the process, we have developed a set of guidelines coupled with working procezdures. Presented in Appendix C, these elements they constitute a guide to content management.

3.2 Workflow for Multilingual Content

Once knowledge-content has been submitted for inclusion in the GSSD knowledge-base, it is subject to the process of knowledge-management. This

[6] Users always have the option of going directly on the Internet and search for their topics of interest, using any of the existing search engines. Recall that in Chapters 1 and 2 earlier we highlighted the key features of the GSSD strategy, notably that its coverage of content is focused, exclusive, and quality controlled.

processs governs the flow of content from its initial identification and submission all the way to its final inclusion and multilingual represetnation in the knowledge-base. The entire process consists of a methodology that has been developed specifically for meeting the GSSD objectives. Since knowledge provision, management, and sharing is a *distributed process,* we must take into account the implications of multilingual and multicultural realities, with all the attendant diversity decentralization.

Initially, GSSD was designed as a knowledge e-networking application focusing on sustainable development and intended to foster the distribution and provision of content, in local contexts, multilingual formats, and distributed geographic locations all over the world. However, a wide range of technological, organizational, and institutional challenges have led to a substantial redesign of the system. As a result, the development of an innovative global workflow application has become a central feature of knowledge e-networking. Thus, the new organizational and institutional priorities and the related technical requirements transformed GSSD into a unique technical application for managing global workflow processes.

3.2.1 Multilingual Workflow Process

When content originates from a single language and location, it is routed through the GSSD workflow-process toward its destination via global workflow until it is converted into all supported languages and resides on all mirror sites of the GSSD network. Clearly, the effective operation of any global knowledge e-networking system depends upon an efficient and effective workflow among the participating actors. In our case, we need to enable users to add content to the knowledge-base, which can then be retrieved by GSSD users.

The concept of *knowledge e-networking* is founded upon collaboration among knowledge-producing actors who retain their autonomy in the course of participating in the e-network, and that this autonomy is a source of value to the actors as well as to the entire e-network itself. Several computational challenges are associated with the effort to support the autonomy of the actors.

Once the content is published to the web, it becomes an extractable part of the GSSD knowledge-base, and is available through GSSD functionalities. This means that a certain degree of intellectual discipline has been exercised regarding the quality of the content and that the submission process has met the basic system requirements. See Appendix C for details. Against this background, Figure 3.1 provides a highly stylistic view of the multilingual workflow process.

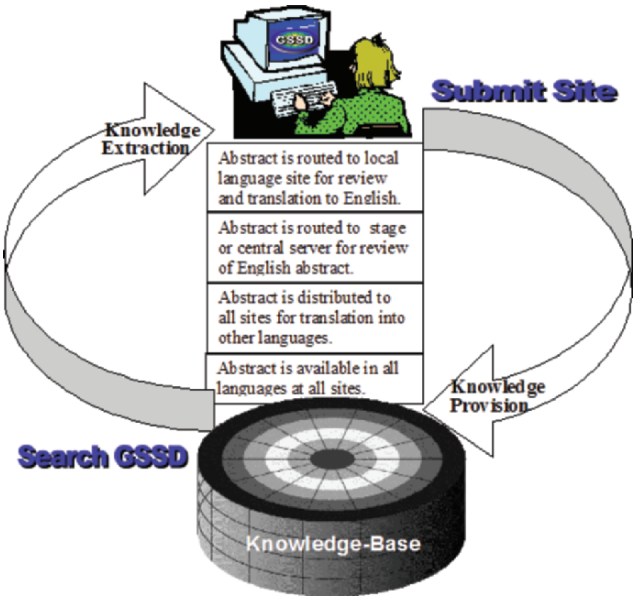

Figure 3.1 **Workflow process: a stylized view.**

For illustrative purposes, we stress the importance of the *original language* of the knowledge-item in question, as well as the *location* of the submission site in the e-network. These are key identifiers that, in turn, determine the precise routing of the submission through the workflow process to its final inclusion in the GSSD knowledge-base.

3.2.2 Routing Knowledge-Content

A more detailed representation of the workflow process is in as signaled in Figure 3.2 which follows the same general sequence of steps noted in Figure 3.1. Consisting of four main segments, the workflow process proceeds as follows: for content–abstract consists of four main steps. Each of these steps is noted in Figure 3.2 along with a representation of the routing process. The steps are as follows:

(1) An abstract–content is first routed to its *local language* site for review and translation into English. For example, if a Chinese abstract were routed to the French server, it would be routed to the Chinese server for translation.

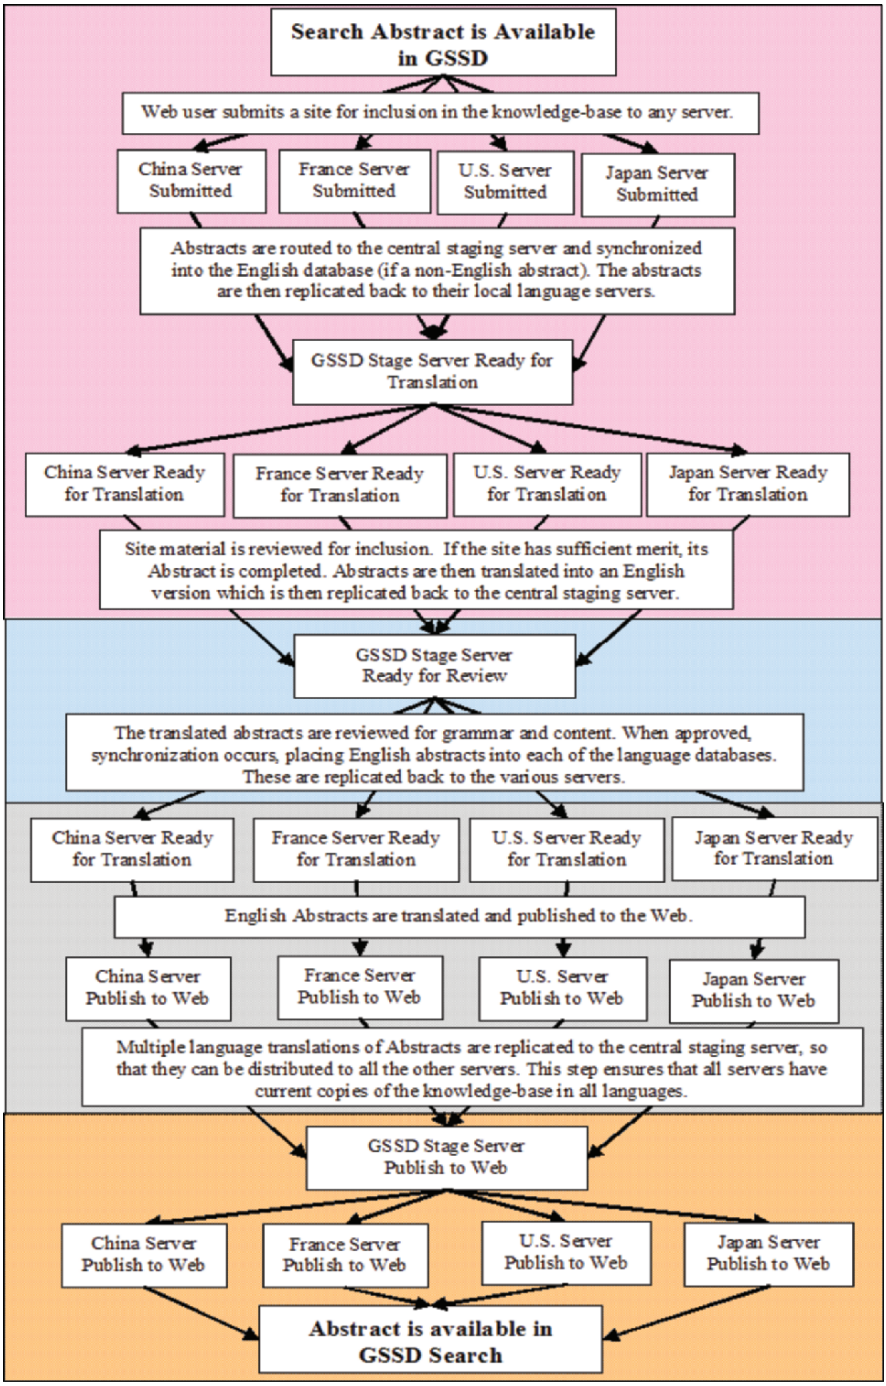

Figure 3.2 **Routing knowledge content: four steps of the global workflow process, differentiated by horizontal lines.**

(2) The abstract is then sent to the central staging server, known as *GSSD Stage*, for review of the abstract in English.
(3) Once the English abstract is reviewed for accuracy, the GSSD Stage server *distributes* the English version of the abstract to all other language sites.
(4) The abstract–content is translated into all *other languages* supported by the system.

When the workflow is completed, the abstract of the new submission has gone through the entire system. Its content is now available in all GSSD-supported languages.

In the discussion so far, we have referred to a general case and to the core principles of the workflow process. However, the strategy allows us to customize the workflow at each step, to meet specific contingencies, while at the same time retaining the integrity of the workflow. We now turn to an illustration of the customization process by comparing four different cases and noting the implications for specific steps in the routing of a submission from its initial entry point until it is integrated in the GSSD knowledge-base.

3.3 Customized Workflow: Four Cases

Given that the workflow is sensitive to the *language of origin*, a certain degree of customization is built into the design of the workflow. At the same time, however, customization does not violate the basic workflow principles for rendering and retaining congruence across all languages in that it is different for English-originated content versus content that is initially submitted in another language.

For illustrative purposes we present a comparative view of four cases with different routing processes that correspond to different situations on-the-ground. Recall that, as noted earlier, the common locale of conversion for consistency is the GSSD Stage Server, and that Stage constitutes the focal point for ensuring congruence across languages and locations.

3.3.1 Case I: English Content Submission

In the first case, the abstract is in English as its point of origin. As Figure 3.3 shows, the workflow process clearly indicates a highly streamlined sequence.

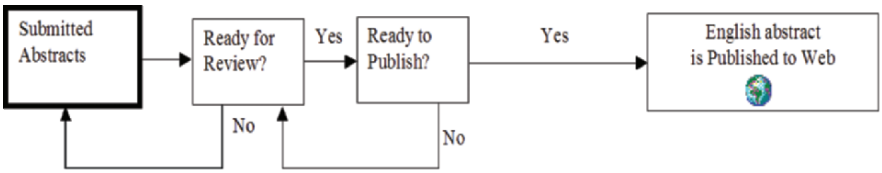

Figure 3.3 **English content-submission.**

3.3.2 Case II: Content Submission Other-than-English

In the second case, the initial content submitted is not in English. Since most of the world's population is not English speaking, providing a process for managing the flow of what will invariably become an increasingly growing number of new participants in e-systems of worldwide. Figure 3.4 shows the routing for Other-than-English submission.

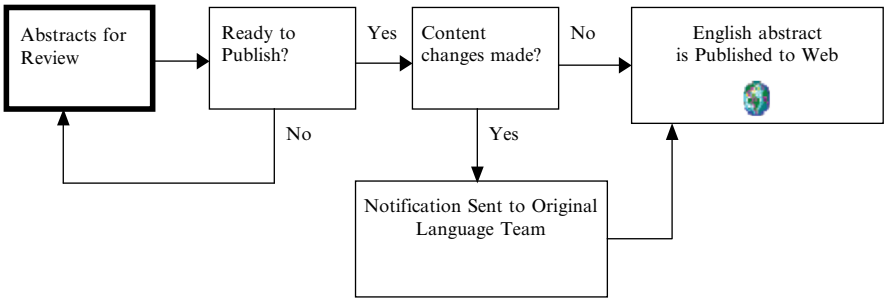

Figure 3.4 **Other-than-English content-submission.**

3.3.3 Case III: Non-English Original Requires English Version

This is a case where the original submission is in Chinese or Arabic, or any other-than-English language, and the contents need to be represented in English (Figure 3.5). This requires the Translation of *** Abstracts into English in order to meet the Stage requirements, and of course, also in order to be placed in the English database.[7] They are also technologically challenging in the representation of meaning, since neither of these are a based on the western alphabet.[8]

[7] The asterisks in this text, namely, *** can represent any non-English language.

[8] Recall that in Arabic, for example, a specific letter may look different if it were located at the beginning, the middle or the end of a word. This issue is fully understood among linguists, teachers, and educators, but is especially daunting for representational systems designed for computational purposes.

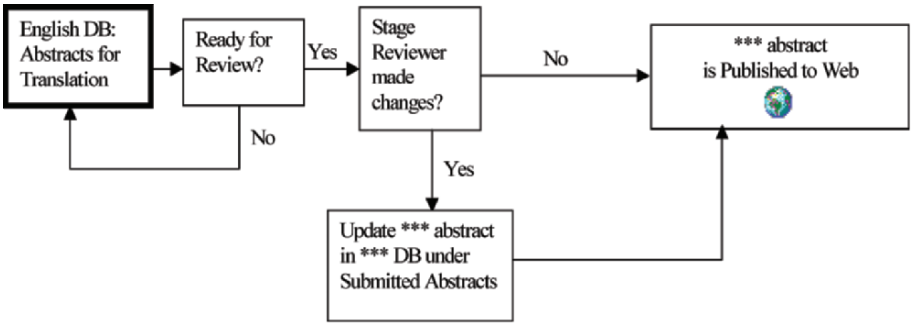

Figure 3.5 **Non-English submission requires English version.**

3.3.4 Case IV: English Submission Requires Other-than-English Version

In this last case, the original content-submission in English needs to be rendered and then published in the non-English languages, including, but not limited to, Chinese and Arabic. The workflow for this case is shown in Figure 3.6.

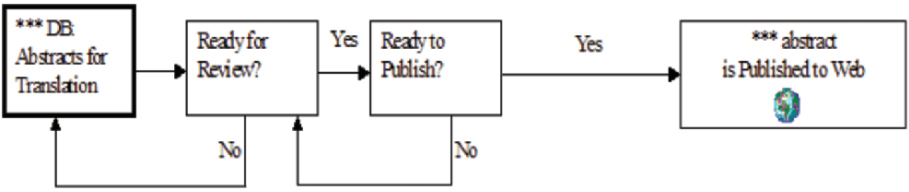

Figure 3.6 **English submission requires other-than-English versions.**

A reminder here is in order, namely that partner in the entire database system (which includes all of the individual language databases as well as the Stage database) must retain its mirror status at every point in time. At first glance, the differences among these four cases may not stand out, but by observing the initial condition at the onset (on the left in each of the above figures) and tracing these to other end of the workflow process (at the right side of the figures) the differences become more evident.

3.3.5 Convergence Glossary

One of the most persistent challenges is to assure the accurate rendering of meaning across languages. This challenge is especially daunting in cases where the sustainable development vocabulary is not fully developed. In such cases,

the Convergence Glossary required for accurate content representation must be developed.[9] This development is essential in order to implement the workflow. The Convergence Glossary consists of the computational representation as distinct from the GSSD Glossary.

The Convergence Glossary is thus is a fundamental feature of the entire GSSD approach to knowledge generation, sharing, management and distribution. The integrity of the workflow process is critical to the robustness of the system as a whole, given the distributed nature of knowledge provision as well as the collaborative management strategy.

As of this writing, there are no precedents for multilingual knowledge provision in non-western languages concurrently for management and sharing on a distributed basis globally. There are few precedents, if any, for the use of western and non-western languages in one integrated and synchronized knowledge-base. Finally, there are limited precedents, at most, for the sustained provision of local knowledge into global networks.

Currently, the knowledge-base consists of Arabic, Chinese, and English. The GSSD team is in the process of updating the French version, and expanding the system and the databases to Japanese and Spanish. Given that the core strength of GSSD lies in its ability to develop and maintain a multilingual and distributed knowledge-base devoted to sustainability issues, extending the language coverage is a key priority.

3.4 Computation and Implementation

So far, we have addressed the workflow process given the diversity in the language of origin, namely, *what* must be done when the submissions are in English versus non-English, and *when* should they be rendered into different languages. We now turn to the question of *how,* namely the computational process by which the GSSD system sustains a knowledge-base that is multilingual in content and distributed in nature.

At this point, we turn to the key operational steps, the logic, and the sequence for developing and managing the design of the GSSD knowledge-base. These issues are important, first, in order to share with others as clearly as possible the methods that we have developed and implemented, and second, to provide foundations on a road map for future efforts that can improve upon, and enhance, these new practices.

[9] The Convergence Glossary is also called the Domino Global Workbench (DGW).

3.4.1 Globalization and Localization

The design of GSSD as a multilingual web application involves turning *one* web application, written in *one* language, into *many* copies of the application in *different* languages – but all linked together conceptually and coherently. The goal is to retain consistency of content and meaning for both knowledge provider and retriever. While the overall GSSD enterprise is governed by the realization of, and the quest for, knowledge *globalization*, this process of consistency-seeking is one of knowledge *localization.* Three sets of e-capabilities, and attendant e-functionalities, are central to the localization process.[10] To be effective these must follow protocol and procedure.

First is application of the *Domino Global Workbench* (DGW), a tool that enables us to automatically create an application in different languages, provided that we first generate a Convergence (or DGW) Glossary of the translated terms. The Glossary serves as a bridge between the language of the original application and the new version created by the process. This step is especially important in the domain of sustainable development where the core concepts are not always available in all languages and a certain degree of conceptual articulation is often needed. In other words, this step involves the development of the vocabulary for managing the knowledge content.

Second is deployment of the *DGW Glossary* which consists of the database that holds all the terms (all text and graphics) that make up the design of the application.[11] Once this Convergence Glossary is created, researchers can then translate the terms in it quickly and easily, with only minimal training on Lotus Notes. DGW then uses this Convergence Glossary to create the translated databases.

Third is *replication* and *synchronization*, two computational steps that render cross-site consistency of content and retains the GSSD mirror sites as identical versions of each other.

These three sets of computational functions are required to keep the application consistent across languages, across servers, and across mirror sites.[12] These features are central to the entire process of multilingual knowledge creation, management, and e-networking. Consisting of both human–computer interactions as well as human–human coordination, they constitute the core mechanisms for distributed global knowledge e-networking.

---

[10] Recall that this process is implemented via the Lotus Domino and related software.

[11] Note that the DGW Glossary (or the Convergence Glossary) is different from the GSSD Glossary. The former is a computational functionality; the latter is about displaying equivalence of meaning, as well as all other concept or vocabulary related issues.

[12] The tasks are usually the responsibility of the GSSD system administrator.

3.4.1.1 Creating a New Language

The DGW tool helps us to manage the localization (i.e. translation) of the GSSD system and render the knowledge-base into local languages and idioms. DGW automatically extracts terminology from the GSSD system and stores it in a Glossary ready for translation. Localized versions of GSSD are built automatically using the translated Glossary. The latter is required to create and maintain consistency of meaning across languages. Also, to accommodate any future changes in the GSSD design, the DGW 'update features' transmit changes easily to the localized versions of content.

3.4.1.2 Working with DGW Glossary

The DGW Glossary is a database that contains all terms in the GSSD design, in the originating language, English, and all translations of terms into supported languages. Each term in the Glossary, in each language, is a separate document. If 12,000 terms being translated into four languages, then 48,000 documents are in the Glossary. Figure 3.7 shows a screen view of the Convergence Glossary, largely for recognition purposes.

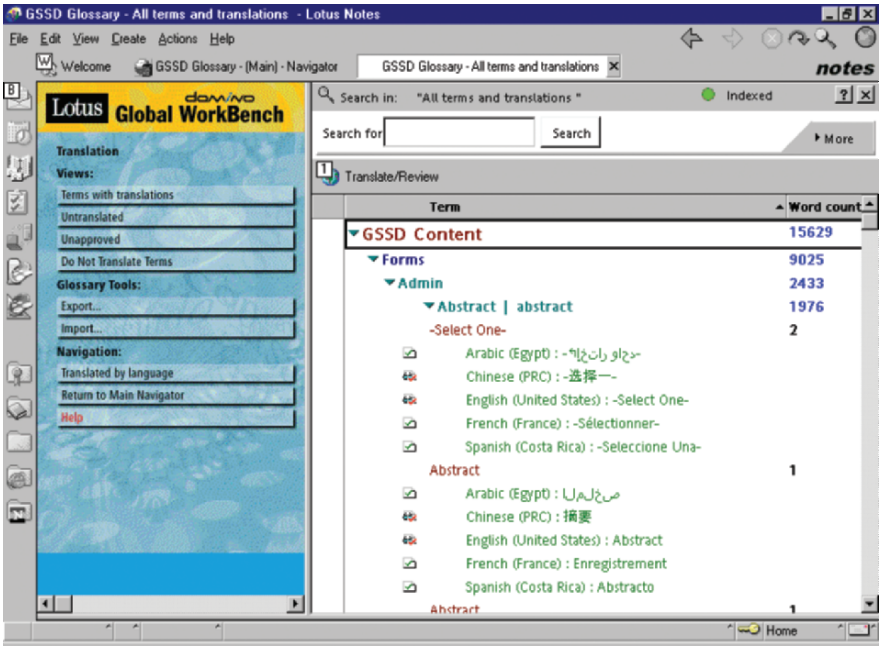

Figure 3.7 **Screen view of convergence (DGW Glossary).**

The DWB Glossary provides another key function. It forces the partners to be as explicit as possible about their understanding of key terms – especially in a domain that is rapidly changing as that of sustainability. In essence, serving as the computational mechanisms for eliminating language barriers, the Glossary provides the groundwork for developing a universal set of multilingual ontologies on sustainability issues.

3.4.2 Replication and Synchronization

As noted earlier, replication and synchronization are two computational processes that fundamental to the operation and maintenance of a multilingual distributed knowledge e-networking system. These processes are usually managed by the GSSD system administrator and constitute the mechanisms by which all sites remain identical in content and languages regardless of location.

Replication refers to the process that transfers changes made in one replica of a database to another replica, thereby keeping the databases located on different servers consistent with each other. A replica is a copy of a database configured in a way that enables replication in the Lotus system. The replication model is a hub and spoke system. The MIT Staging Server acts as the hub. It is not live to the web and functions as the central, up-to-date repository of the application. The spokes are the live mirror sites, which contain replicas of the application. The Staging Server replicates with these servers on a scheduled basis, to ensure that all mirror sites have consistent and updated content.

Viewed in these terms, replication is essentially a process of transformation-for-consistency. During the process of replication, the staging server at MIT (GSSD Stage) interacts with each mirror site to identify what has changed at the Live Site and on GSSD Stage, and then sends or receives those changes as relevant.[13] At the end of the replication process, the GSSD application is identical on each server. For example, a user looking at the application on the Chinese server will see the same content as a user looking at the MIT server. This is the point at which processes of globalization and localization converge, in theory and in practice.

GSSD Stage regularly replicates with each Live Server. This means that any change made in any server anywhere is then automatically rendered also on every other site. Any changes, introduced in any of the system-locations propagate through the entire GSSD network. For this reason, each computational

[13] Changes include modifications, additions, as well as deletions. For example, if an abstract in one replica is deleted from the knowledge-base, then during replication, all other replicas of that abstract are also deleted.

feature as well as the overall replication process requires careful implementation.

Synchronization refers to the process that through which consistency of content across languages is maintained. When a document is created and reviewed in the English language database, and the synchronization process is undertaken, a copy of that document is created in all other language databases. This function enables human translators render the content from English into their local languages. This is the only point at which human translation takes place. Since the translation is specifically of knowledge-abstracts, the resort to human skills is focused as well as limited. The GSSD system uses English as the reference language, which means that all content in GSSD needs to be in English system-database before it can be processed into other languages.[14]

Once a document is synchronized, it becomes virtually linked to its counterparts in all language databases. In other words, the operating system identifies it as the *same* content document located in multiple language databases. For example, if a synchronized document is deleted in its originating language database, then the corresponding documents in the other language databases will be automatically deleted during the next synchronization process.

Finally, synchronization is undertaken routinely on a regular schedule on the GSSD Staging Server at MIT. Any changes due to synchronization on GSSD Stage will be replicated to the Live Servers in the next replication process.

3.5 Conclusion

This chapter focuses on the challenges of creating and managing multilingual knoweldge content. It demonstrates the transition from a basic knowledge e-system to global knowledge e-netowrk via application of the workflow process. Most of the issues addressed so far are methodological in nature, supported by screen views for the purposes of illustration. At the same time, however, this chapter represents the operational manifestation of some of the concepts and theory underpinning sustainability presented in Chapters 1 and 2.

The 'Guide to Core Concepts' in Appendix A serves as a link between the theoretical segments of this overall initiative at the *conceptual* plane, on the one hand, and the methdological features at the *implementation* level, on

[14] Note that the English language database is distinct from the English system database. The former is comparable to any other language database. The latter serves as the convergence needs. The fact that they are both in English should not obscure the difference in the function within the overall workflow process.

the other. While seemingly only technical and procedural in nature, the issues addressed in this chapter are at the very heart of a distributed knowledge system. They pertain to the selection of knowledge, to its abstracting, to the practical application of conceptual guidelines, and to the indexing of content.

The value of knowledge lies in its content. In the absence of robust methods for valuing knowledge content, it is incumbent upon us to articulate our operational rules for knowledge provision. Increasingly, the international community as a whole, and the scientific, educational and knowledge community more specifically, are recognizing the imperatives of decentralization in knowledge provision and access and the importance of a concerted 'push' toward localization.

From the user's perspective the details of the workflow – notably replication and synchronization – need not be examined in any great detail. But for a system manager, mirror-site partner, or other partner types (noted in Chapter 4) the details have operational implications.

More fundamentally, this chapter shows some of the practical, strategic, and essential tasks that accompany any collaborative knowledge-based initiative – particularly those of global scale and scope – and to transform this collaboration into its multilingual e-rendering. In short, the challenge is to establish and maintain the capacity to engage in knowledge provision, as well as knowledge retrieval, over an integrated knowledge-base originating from diverse, distributed, and decentralized sources.