

United States Air Force Scientific Advisory Board



Report on Building the Joint Battlespace Infosphere Volume 1: Summary

SAB-TR-99-02

December 17, 1999

Cleared for Open Publication – February 2000

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 17 DEC 1999		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Report on Building the Joint Battlespace Infosphere Volume 1: Summary				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of the Air Force AF/SB Washington, DC 20330-1180				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 188	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

*This report is a product of the U.S. Air Force Scientific Advisory Board Committee on **Building the Joint Battlespace Infosphere**. Statements, opinions, recommendations, and conclusions contained in this report are those of the Committee and do not necessarily represent the official position of the U.S. Air Force or the Department of Defense.*

United States Air Force Scientific Advisory Board



Report on Building the Joint Battlespace Infosphere Volume 1: Summary

SAB-TR-99-02

December 17, 2000

Cleared for Open Publication – February 2000

(This Page Intentionally Blank)

Executive Summary

The Joint Battlespace Infosphere Defined Operationally

The Joint Battlespace Infosphere (JBI) is a combat information management system that provides individual users with the specific information required for their functional responsibilities during crisis or conflict. The JBI integrates data from a wide variety of sources, aggregates this information, and distributes the information in the appropriate form and level of detail to users at all echelons. The JBI was originally described in the 1998 USAF Scientific Advisory Board (SAB) report *Information Management to Support the Warrior*.

At the joint task force (JTF) commander's level, the JBI is a powerful command and control (C²) system that combines inputs from a variety of sources, including existing C² systems, reconnaissance data, satellite data, unit capability data, logistics data, and real-time battlefield conditions. The JBI builds an aggregated picture from these combined inputs, giving unparalleled situational awareness accessed as easily as a web page. The JBI also provides for speedy downward flow of information, so when commanders order an action, the action is received and implemented at the subordinate level almost immediately.

The commander in chief (CINC) or JTF commander creates a JBI for a specific purpose, usually in response to a crisis or conflict. The JBI enables the commander to focus information support for a specific operational purpose, ensure or limit access to critical information, and provide an information management system that can respond to natural or enemy actions that disrupt communications capabilities. As units are assigned to the mission, their information needs are electronically identified, and available information is automatically accessed. Thus, deployed units are ready to fight immediately upon being deployed or assigned.

The Joint Battlespace Infosphere Defined Technically

Supporting these capabilities and forming a foundation of the JBI is a platform of protocols, processes, and common core functions that permit participating applications and organizations to share and exchange critical mission information in a timely manner. It provides uniform rules for publishing new and updated objects into the JBI and promptly alerts any JBI clients that have subscribed to such objects. These properties enable dynamic information flows among client programs of the JBI, serving to integrate the clients to conduct a single mission.

The JBI platform integrates many individual information systems that currently support operational forces. Each existing system has been developed in a stovepiped fashion; few interoperate with each other. The JBI acts as an intermediary between these systems, converting information from one representation to another to enable interoperability. In addition to acting as middleman between disparate systems, the JBI interprets the information flowing between applications, using it to build its own, more complete, picture of the current situation. Furthermore, the JBI tailors this picture for individual users: the commander gets a high-level view of the campaign, while the soldier in the field gets a detailed description of a nearby hostile base.

The JBI provides an architecture for the incorporation of future data capture technologies that exploit better sensors, databases, fusion engines, automated analysis tools, collaborative planning and execution aides, and distribution controls. It is also a disciplined process that guides the activities of people responsible for obtaining, verifying, fusing, presenting, analyzing, and controlling the information necessary for success in any operation.

The Joint Battlespace Infosphere Defined Relative to Present Systems

The JBI is connected to, and interoperable with, a variety of existing and planned C² and combat support information systems. The JBI is not intended to replace C² systems, but to be the substrate for integrating them. The JBI subscribes to pertinent information published by supporting systems and, when necessary, pulls specific information from other networks. In addition, the JBI connects to fusion engines and may perform fusion on its own, thereby ensuring that the most complete and coherent picture of the battlefield situation resides within the JBI itself. The JBI concept recognizes that display technology is constantly advancing and that new displays must be tailored for users from flight leader to JTF commander.

The JBI provides services through a federation of multiple servers. The Global Information Grid connects these servers to each other and to the many systems that support the JBI. Many of the servers provide services from the rear via reachback, thereby limiting the forward footprint of the JBI.




Figure S-1. The JBI integrates C² resources on top of the Global Information Grid infrastructure

Report Overview

This report provides a brief description of the JBI's capabilities and its operation. The main focus of this report, however, is the technical roadmap for building the JBI. Namely, it answers the questions "What technologies make up the building blocks of the JBI?" "In which technologies should the Department of Defense invest?" and "How should that investment be managed?"

JBI Capabilities

What does the JBI do for the warfighter? First, *warfighter* is defined in the broadest sense, from intelligence analyst to maintenance crewmember to JTF commander. Each warfighter receives exactly the information needed to perform his or her function. Not only is it the right information; it is updated when new information enters the JBI, so that the information delivered is timely and consistent with the information shared with all JBI users. Furthermore, that information is delivered via an appropriate interface: on a soldier's handheld computer, the pilot's head-up display, or the virtual-reality collaborative environment within which many users share information and explore alternative courses of action through simulation.

Not only does the JBI distribute information, it aggregates and fuses information to generate higher-level knowledge. Thus, the joint force commander (JFC) can get an aggregated, or summary, view of the battlespace. The commander can also request more detailed information on a particular topic of concern and the predicted consequences of command decisions. The JBI maintains pedigree information as it aggregates information, so the commander can drill down to examine the inputs and processes the JBI uses to generate an aggregated piece of higher-level knowledge.

In addition to providing information to users at all echelons, the JBI eases system management through automation of critical tasks. The first of these tasks is to stand up the JBI appropriate for the region and mission to be performed. A critical function is the management of units as they are assigned to the JTF. Each deployed unit defines its capabilities, support needs, and information interface (subscription and publish descriptions) through the use of force templates. Force templates define the electronic handshake between the JBI and subordinate units, and their use lets units be quickly added to the JBI with little or no manual reconfiguration required. Another self-management task performed by the JBI is bandwidth management. Specifically, the JBI must ensure that communications links are used efficiently so that information is rerouted or volume reduced when links are down, degraded, overloaded, or compromised.

The JBI also supports management by the JFC's information management staff. This staff can change access controls on information, write scripts (that is, fuselets) to redirect or aggregate information in a new way and to configure gateways between the JBI and coalition information resources.

Because the information staff can change the JBI by writing new fuselets, the JBI is flexible. The JBI serves the JFC and other users in performing their jobs rather than constraining them to perform in fixed, awkward ways. Furthermore, operational processes for working with the JBI

will develop as the JBI is in spiral development, so any system tweaking and tuning needed in a deployed JBI will be minimal.

In sum, the JBI enhances information storage and information flows among the people and computer processes engaged in conducting a military operation. Improved ability to sift and distill information rapidly provides better guidance to the commander, staff, and warfighters of a mission. A mission is configured with the right information-processing resources, both humans and computers, to manage the mission's information and develop useful knowledge from the information. The JBI's role is to store or provide access to sensor information, intermediate results, and ultimate knowledge in a repository so that it can be shared throughout the mission—subject to proper access authorization. The JBI also arranges to route information to the right destinations, alerting the people and processes that should respond to new data. The chain of alerts constitutes a workflow process, designed and adapted to process the mission's information. These JBI mechanisms ensure that the information it provides is an asset to the mission.

JBI Technical Overview

The JBI is built on four key concepts. These are

1. Information exchange through “publish and subscribe”
2. Transforming data into knowledge via fuselets
3. Distributed collaboration through shared, updateable knowledge objects
4. Assigned unit incorporation via force templates

The next sections describe the use of these technologies in the JBI implementation.

Information Exchange Through Publish and Subscribe

Users and programs *subscribe* to information of interest in the JBI. Each piece of information is stored in the JBI as an *object*. Objects contain both the represented information and metadata (that is, data about data) describing the information. Subscriptions contain search values for metadata fields. When a new object is published in the JBI, any subscriptions matched by the object's metadata are fulfilled. That is, the subscriber receives the new information. If the subscriber is a user, the user receives an immediate notification of the new information. The form of the new information depends on the user, the application program, and the computing device being used. For example, the JBI may overlay a new graphic on a 2-D map image, or change the list of weapons to be loaded on an F/A-18, or sound an audible alarm in conjunction with flashing red visuals on a cathode ray tube. The next section describes the situation when a program subscribes to information.

Transforming Data Into Knowledge Via Fuselets

Programs subscribing to information in the JBI are called *fuselets*. When information becomes available to a fuselet via a subscription, the fuselet publishes one or more new objects as a result. Fuselets can be used to encode a commander's standing orders, such as “if a tactical ballistic missile launch is detected, issue a ‘major threat’ alert.” Fuselets may also be used to provide regular reports that summarize information, for example, weapon inventories and sortie counts. The JBI includes a library of fuselets to cover anticipated situations, and the information

management staff may write more fuselets using a scripting language. While fuselets can act on certain conditions, they are based on rules and do not themselves exhibit common sense or judgment when executing. However, fuselets may submit requests to more robust and complicated fusion engines that apply reasoning to the input data. Eventually users may implement fuselets for specific purposes by dragging and dropping an icon to carry out specific functions.

Distributed Collaboration Through Shared, Updateable Knowledge Objects

Users interact with the JBI in many ways. Behind these interaction modes is a set of information objects describing the common operating picture. When new data or information arrive in the JBI, subscribers receive the appropriate direct information or as derived through the use of fuselets. For example, a command center user may have a 3-D virtual reality environment modeling the entire battlespace. A soldier in the field may have a personal digital assistant. If the soldier in the field reports the presence of an artillery battery, the user in the command center may see the battery instantly in the virtual reality model. Similarly, the JBI supports collaborative planning using the “shared whiteboard” notion—that is, collaborative tools let multiple users interact with an application, see changes made by other users, and ultimately come to agreement on the final product.

Unit Incorporation Via Force Templates

The force template is a software description of a military unit that may be integrated into a JTF. Several varieties of force templates are used. One form of force template describes a fighting unit (for example, a tank battalion or fighter squadron). The key elements of information included in such a force template include force employment capability, ammunition inventory, fuel requirements, communications requirements, computing systems, information requirements (the unit’s clients’ subscriptions), information products (objects to publish both during instantiation and later)—for example, intelligence, surveillance, and reconnaissance (ISR), and personnel requirements. Another variety of force template is one describing a support unit. A template for these units always includes information requirements (the unit’s clients’ subscriptions), information products (objects to publish), communications requirements, and computing systems. However, the specific publish and subscribe exchange described in the force template varies depending on the type of unit. For example, an airlift control unit force template would agree to publish information about the locations and movement of relevant airlift assets while subscribing to airlift needs generated by the JBI.

Technical Infrastructure

Other technologies support the JBI as well. These are briefly described in this section.

Browsing. JBI users browse through the JBI, much the way they browse the Web today. The browser is able to fetch objects from the JBI, to search the JBI using queries, and to make visual presentations of many of the common JBI information objects. A browser may also include tools for creating objects, which then are published to the JBI.

Interaction. Many clients are designed to connect humans to JBI information in special ways. For example, a mission rehearsal client might query the JBI to obtain details of a mission and then build a 3-D fly-through environment in which a pilot can rehearse the mission.

Fusion. One of the missions of the JBI is to fuse information from a variety of sources into high-level “knowledge” that is readily accessible to the commander and other staff. Fusion programs designed to take advantage of the JBI’s information exchange capabilities can subscribe to information from many different sources. If the information is available, the fusion program is informed via the subscription, and the program can fuse the information. In other words, fusion programs can now be written to take advantage of any available information they can understand and evaluate.

Objects. The JBI information objects are like extensible markup language (XML™) documents, instead of objects in the sense of object-oriented programming. Every JBI object is an instance of an object schema. The object schema defines and abstracts a category of things using <attribute, value> pairs. Object schemas are stored in the JBI. Client programs are able to query the JBI to discover the existence of an object class that may satisfy their information needs. Object metadata attributes describe the JBI object rather than the real-world object represented by the object. The metadata values support querying and subscriptions.

Structured common representation. Objects in the JBI are related to each other. The objects and their relationships form a representation of the current military situation. These object relationships are stored in a structured common representation (SCR), which describe hierarchical relationships as well as more *ad hoc* relationships. The SCR supports presentation and tailoring of information. It also supports drill down through hierarchies of knowledge, so the user is able to examine evidence supporting presented information.

Automatic data capture. As users interact with the JBI, the interface used is natural for their duties, locations, and types of devices. Voice recognition technology is an obvious step from keyboarding to a more natural interface. As technology matures, the interface could infer information from both the user’s voice and gestures, so that when a user points at a display and says, “There,” the system could infer the intended location.

Tailoring information to meet user needs. The understanding of a situation or the available options depends critically on presenting the information in an appropriate form. The presentation format exploits multimodal sensor input from a combination of visual, aural, and haptic interfaces. Geospatially referenced presentations, 3-D graphics, animation, image zoom, and moving forward and backward in time are some of the tools that are available. But the presentations must be tailored to the workflow task and to the preferences of a particular user. What is presented in the cockpit may be very different from what is presented in a command center.

Implementation Possibilities

The JBI implementation will not result in a monolithic, single-paradigm program. Instead, it will use different approaches to provide different services. The most promising candidates for the JBI architecture from today's technologies are listed below.

Digital libraries. A digital library consists of multiple repositories responsible for storing digital objects. Digital libraries include *index servers* to answer queries and searches, while *handle servers* provide a cross-reference between objects and the repositories in which they reside.

Enterprise integration technologies. This is a class of middleware techniques used to integrate otherwise separate (that is, stovepiped) business information technology applications. Generally, they use software "connectors" to provide a linkage between each application and a shared structure for communicating information from one application to another.

XML technologies. XML is a hypertext markup language (HTML) successor. XML provides a way to describe data structures in a textual form. It can be used to represent JBI object schemas and objects. Furthermore, commercial interest in XML has led to commercial-off-the-shelf (COTS) XML tools that are proliferating even now.

A Development Roadmap

The core recommendations resulting from this study are

- Immediately fund concept validation prototypes. Each partially complete prototype (YJBI) will employ an object-based common representation. Each will include an object repository and object-based force templates. Low-cost experimental prototypes will demonstrate a service, such as publishing information from a YJBI client¹ to the object repository. YJBI experiments should make heavy use of commercial products, even if these may include products inappropriate for later operational designs for reasons of security or other factors.
- Initiate a information management cadre to work operational business processes as part of the YJBI development.
- Support or reallocate funding for research programs within Service laboratories and the Defense Advanced Research Projects Agency (DARPA) to support advanced JBI platform concepts. These include issues such as common representation, military information assurance, and information fusion.
- Initiate spiral development of an operational JBI. Initiation should occur after designers and warfighters agree on functionality, and after YJBI concept validations have established credibility.

The study team recommends that the concept validation and early spiral development be based heavily on web technology. In the Department of Defense (DoD), developmental vectors for many C² systems are realigning toward the web. In addition, other commercial products must be evaluated for possible deployment in the JBI, and they should be included as early as possible during spiral development.

¹ A JBI client is any software application that makes use of JBI platform services to publish, subscribe, or otherwise interact with JBI information objects.

- Development of the long-term JBI architecture should not impede rapid development of prototypes, since early prototypes are crucial to the warfighter and provide architectural validations. Similarly, the youthful state of the art in information assurance should not slow the development of the JBI. DARPA and others have large research efforts in information assurance. Rapid prototypes need only exhibit partial coverage of JBI core platform services: a skeletal common representation, at least one C² client application exhibiting publish and subscribe, and an interaction capability. Rapid prototypes should be designed to clarify the vision of USAF warfighters, joint-Service users, and C² system design specialists. The focus of these first JBI prototypes should be firmly on *inexpensive evaluation and idea generation*. Architecture studies conducted in parallel should focus on downstream functionality for spiral development. While early prototypes provide early concept validation and feedback, long-term research must be supported to provide critical capabilities in later stages of spiral development.

Finally, the spiral development of the JBI must go hand in hand with development of the information staff, both professionally and technically, and with the development of the JBI business processes. The C² business processes are tightly linked to the information systems in use. These business processes must evolve in spirals with the JBI. In addition, new business processes should be evaluated using process models, data models, workflow models, and simulation models.

The JBI Roadmap

The JBI study team created a roadmap for development and phased improvement in the format of major weapon systems. The roadmap serves as a programming guideline and is summarized below.




Figure S-2. JBI development roadmap

Executable Recommendations

To ensure that the JBI goes from the vision described herein to reality, the JBI must be considered a major weapon system. In that light, this study makes the following additional recommendations relating to the management of the JBI:

- Create an information staff function
- Develop new concepts of operations at the Aerospace Command and Control Intelligence, Surveillance, and Reconnaissance Center (AC²ISRC)
- Define common information representations led by the Assistant Secretary of Defense for Command, Control, Communications, and Intelligence (ASD-C³I)
- Reinforce DARPA research and development (R&D) investment for JBI technologies
- Focus the Air Force Research Laboratory (AFRL), other Service research labs, and battlelabs on evaluating and applying commercial technologies for the JBI
- Create the JBI testbed now for Joint Expeditionary Force Experiment 00 participation
- Link the JBI testbed to other Service efforts in digitized battlefield and network-centric warfare
- Promote the JBI to the CINCs

(This Page Intentionally Blank)

Table of Contents

Executive Summary	iii
List of Figures	xiv
List of Tables	xiv
Chapter 1: The Joint Battlespace Infosphere	1
Chapter 2: An Operator's View of the Joint Battlespace Infosphere.....	13
Chapter 3: Operating the JBI	27
Chapter 4: Warfighters Using the JBI.....	37
Chapter 5: The JBI Input Segment.....	63
Chapter 6: The JBI Platform	75
Chapter 7: Technologies for Building the JBI	101
Chapter 8: Building the JBI	133
Chapter 9: Overall Recommendations	147
Acronyms and Abbreviations	151
Appendix A: Terms of Reference	157
Appendix B: Study Members and Affiliations	159
Appendix C: Study Meetings.....	161
Appendix D: Spiral Development.....	163
Report Distribution	167

List of Figures

Figure S-1. <i>The JBI integrates C² resources on top of the Global Information Grid infrastructure</i>	iv
Figure S-2. <i>JB I development roadmap</i>	xi
Figure 1. <i>The Battlespace Infosphere functions described in the 1998 study: Input, Manipulate, and Interact</i>	6
Figure 2. <i>The JBI integrates existing and future C², ISR, and support systems</i>	28
Figure 3. <i>The information management staff organization. The staff includes functional expertise from ops, intel, logistics, and communications</i>	29
Figure 4. <i>Functions of the Interact Portion of the JBI</i>	37
Figure 5. <i>Commander subscribes to current status information to monitor the area of responsibility</i>	38
Figure 6. <i>JB I transforms alert to support commander's task and preference</i>	39
Figure 7. <i>Collaborative Problem Solving</i>	43
Figure 8. <i>History of Napoleon's campaign against Russia (five dimensions displayed in a single graphic: location, temperature, size, battle location, name, and date)</i>	46
Figure 9. <i>Automatic formatting and filtering of information for field and cockpit</i>	51
Figure 10. <i>Automatic data capture for combat support</i>	55
Figure 11. <i>The JBI Input Segment deals with inputs such as combat support products, fusion products, planning and execution products, command guidance, and user information products</i>	63
Figure 12. <i>Inputs to the JBI</i>	64
Figure 13. <i>An evolving "information web" will provide a link from the JBIIS to existing legacy and evolving military C² assets</i>	68
Figure 14. <i>Service-based middleware will allow the interoperability of heterogeneous systems and the retrieval of information from legacy systems</i>	70
Figure 15. <i>The JBI Platform provides infrastructure JB I services: publish, subscribe, query, and control</i>	75
Figure 16. <i>Schematic JB I Platform structure</i>	81
Figure 17. <i>Example object: an air tasking order using XML encoding for the object, with an embedded XML-MTF encoding of the ATO</i>	85
Figure 18. <i>Information technology evolution and the dangers of too rigid specifications</i>	104
Figure 19. <i>Architecture of commercially available middleware</i>	105
Figure 20. <i>Emerging Commercial Availability of Component Technology</i>	110
Figure 21. <i>Operational view of dynamic retargeting</i>	142
Figure 22. <i>Information-centric view of dynamic retargeting</i>	142
Figure 23. <i>JB I Development Roadmap</i>	144
Figure 24. <i>JB I metrics</i>	165

List of Tables

Table 1. <i>Example of an Active Template. When the location is entered (in the right-hand column), the generic values change to tailored values for that location.</i>	48
<i>(Question marks identify information requested from the user.)</i>	48
Table 2. <i>Example Metadata Tags.....</i>	87
Table 3. <i>Network Capabilities</i>	115
Table 4. <i>Example of a JB I spiral.....</i>	163

Chapter 1: The Joint Battlespace Infosphere

1.0 Introduction to the Joint Battlespace Infosphere

The Joint Battlespace Infosphere (JBI) is a combat information management system originally described in the 1998 USAF Scientific Advisory Board (SAB) report *Information Management to Support the Warrior*. The JBI integrates information from a wide variety of sources, aggregates the information, and distributes information in the appropriate form and level of detail to users at all echelons. At the joint task force (JTF) commander's level, the JBI is a powerful system that combines inputs from a variety of sources, including command and control (C²) systems, reconnaissance data, satellite data, unit capability data, logistics data, and real-time battlefield conditions.

The JBI performs integration, aggregation, and distribution by combining human control, supported by collaborative decision-making environments, with automated rule-based decisions to achieve information superiority leading to optimal employment of fielded forces. The complex nature of the JBI dictates that data be organized by referencing and cataloging. JBI data can then be published, subscribed to, or searched to meet user needs. Data are continuously created and updated as a result of automated fusion of sensor data, user input, and propagation of information throughout the JBI. The wealth of information in the JBI is aggregated and summarized at the right level of detail for each user.

The JBI is not one centrally controlled system that supports all operations worldwide. Rather, a JBI is established when a commander in chief (CINC) deems it necessary, based on the development of a crisis or contingency. Some JBIs will remain in constant operation to support potential conflicts, as in Korea. A JBI is initialized by the CINC's information management staff to respond to the situation. The CINC or JTF commanders and their functional staffs enter information about the commander's intent, available forces and their capabilities, the forces' maintenance requirements, the forces' information requirements, ISR capabilities and needs, and other campaign-specific information into the initialized JBI. This initialization defines the kinds of information objects to which users and systems may subscribe. As objects are created and updated, subscribers are kept up to date. As new units are deployed, they seamlessly become part of the JBI, subscribing to and publishing relevant information from the start.

With current information at their fingertips, warfighters make more effective decisions more quickly. Changes in weather, the ground war situation, bomb damage assessment, intelligence information, or weapon availability are handled in near-real time. When maintainers publish via the JBI the status of a newly repaired F/A-18, planners then assign the plane a target to attack. When a platoon of ground troops requests the location of enemy tanks, the JBI provides that information in a form tailored for the personal digital assistant carried in the field. At a higher level, the CINC sees an aggregated summary of the entire campaign with an option to "dip a cup" into the JBI to get more details on any aspect. Immediate access to needed data for all JBI users leads to decisive information superiority for U.S. forces.

1.1 Motivation for the *Building the Joint Battlespace Infosphere* Study

Information Management to Support the Warrior defines a Battlespace Infosphere (now referred to as the Joint Battlespace Infosphere) and states that a JBI is needed to manage combat information. Several previous Defense Science Board and Army Science Board reports also describe the need for an information management system similar in many ways to the JBI. The JBI or JBI-like systems described by those studies propose the integration of planning, execution, C², intelligence, and surveillance information systems to provide full, real-time battlespace awareness to users at all echelons.

It is easy to understand why these science board studies unanimously champion the building of a powerful system with capabilities that are significant force multipliers. While providing great vision in applying information technology to its full potential, all the studies lack the technical details the Services need to start development of the JBI.

This report fills the void left by the earlier studies. It describes the existing technological components that should be built into the first JBI prototype over the next 2 years. It contains a concrete description of the architectural foundation for the JBI of the future using promising commercial technologies. This report also assesses candidate technologies to guide the JBI's spiral development as well as R&D investment. In short, this report provides a technical roadmap for building the JBI.

1.1.1 Organization of This Chapter

The next section describes the ways the JBI supports current military initiatives, especially in the realm of information technology deployment. The subsequent section briefly summarizes previous studies to show how their conclusions lead to the development of the JBI. The 1998 USAF SAB study is described in more detail to show the high-level architectural design for which the current study provides an implementation strategy. Following that is the charter for the study and the approach taken by the study members to meet the charter.

1.2 The JBI and Current Military Initiatives

The JBI should not be considered only in the context of studies. Just as important, the JBI enhances current changes in business practices taking place in all the Services. Many of these practices focus on the use of information technology, but the JBI supports other initiatives as well.

The Air Force is evolving into an Expeditionary Air Force. The JBI supports this by being easily and quickly deployed to support all varieties of contingency operations. By providing leading-edge information management and battlespace awareness, the JBI forms a cornerstone on which the Air Expeditionary Force can be effectively deployed. In addition, the JBI rapidly distributes information, rendering the JBI's geographic location unimportant. The result is that reachback can be used for many functions, so the forward footprint of the JBI is small, in keeping with the Air Expeditionary Force philosophy.

The Navy's operational concept, *Forward ... From the Sea*, depends on superior speed of command to enhance the advantages of operating from the sea. Speed of command is the ability to rapidly collect information, assess the situation, develop a course of action, and immediately execute with overwhelming effect.

The Marine Corps' operational concept, *Operational Maneuver From the Sea*, fits within the *Forward ... From the Sea* framework. The Marines expect to operate at an overwhelming tempo, so they require C² systems oriented toward rapid decision-making at all levels of command. Systems which reduce uncertainty, provide battlefield visualization, and select critical information are identified as supporting the *Operational Maneuver From the Sea* concept. Furthermore, C² systems must emphasize intelligence, deception, and flexibility while integrating organic, joint, and combined assets. Combat service support cannot be ignored, as limited resources available to handle and deliver material must be managed to deliver "right time, right place" support in the face of rapidly changing requirements.

The Army's vision for its future is named the "Army After Next" (AAN). One of the key tenets of this vision is that, whatever the strategic situation, whatever the combat mode, the soldier will be utterly dependent on information technology. The "Army After Next" depends on an information management system that provides common situational awareness to friendly forces, real-time intelligence on enemy forces, and fire control. This system relies heavily on artificial intelligence for its management functions and responsive database technology to speed queries. Finally, the Army's information management system must support multilevel security.

As the description of the JBI unfolds in this report, it will become apparent that the JBI's capabilities fit the visions espoused by the future operational concepts of the different Services. Moreover, a joint system integrates the information from all Services so that knowledge and functionality are shared throughout the battlespace, not isolated in stovepiped systems.

Many of the capabilities envisioned in the Services' future operational concepts have germinated from scientific board studies. The next section summarizes several of these studies.

1.3 Previous Studies—A Common Theme

This section briefly describes several previous studies concerning information management and shows how the JBI system described in this study logically follows from the recommendations of the sequence of studies.

1.3.1 1994 Air Force SAB Report: Information Architectures That Enhance Operational Capability in Peacetime and Wartime (SAB TR-94-002)

Information Architectures That Enhance Operational Capability in Peacetime and Wartime does not define an information architecture per se; rather it contains principles to guide the development of large information systems. The more important principles include following commercial trends; use of spiral development; use of open systems with well-defined, extensible, layered services; use of a standard security architecture; use of a common data dictionary to enhance interoperability; enforcement of standards at interfaces rather than in components; and hierarchical control starting at the Department of Defense (DoD) level, where the architecture is

defined using “building codes,” and where “building permits” are used to enforce conformance to standards. This study also recommends that a major command, control, communications, and intelligence (C⁴I) interoperability testbed be established to draw together disparate activities of the joint Services and develop broad experience with existing and proposed commercial standards.

1.3.2 1995 Air Force SAB Report: *New World Vistas*

The “Information Technology” volume of *New World Vistas* provides discussions of technologies in which Air Force R&D resources should be invested. Some of these key technologies are human-computer interaction (HCI), including telepresence and augmented reality; automated information fusion combining both signals and symbolic knowledge; and reasoned-action and learned-action software agents. The “Information Applications” volume describes networking infrastructures and both offensive and defensive information warfare. Also included are recommendations for improving situational awareness using the technology mentioned above. A proposed model for planning and execution is described as a dynamically updated 3-D spreadsheet relying on push and pull paradigms, agent technology, and advanced HCI. Finally, requirements for dynamic C² specify the need for cutting decision time from hours to minutes and for evolving doctrine with new technology.

1.3.3 1996 Air Force SAB Report: *Vision of Aerospace Command and Control for the 21st Century* (SAB TR-96-02)

Vision of Aerospace Command and Control for the 21st Century states that today’s C² systems are not broken, but that they do impose limitations on combat effectiveness. This study recommends that future C² systems include enhanced decision-making tools for solving multidimensional, time-sensitive problems. C² systems must be modular to enable tailoring for specific use with a minimal logistics footprint. They must rely on commercial infrastructures where logical and reliable. They must make information available through integration, interoperability, and tailorable releasability to all operators. Finally, C² systems must provide a “plug and play” capability for quick and effective response to any operations.

1.3.4 1997 Air Force SAB Report: *United States Air Force Expeditionary Forces* (SAB TR-97-01)

More than a fourth of the summary volume of the *United States Air Force Expeditionary Forces* study discusses command, control, and intelligence. Much of that discussion describes technology needed to support Air Expeditionary Forces. The key enablers needed to support an Air Expeditionary Force are listed as global connectivity; information management; battlespace awareness; geospatial position, navigation, and timing; and system assurance. Connectivity is provided by commercial off-the-shelf (COTS) communication services where possible, with reserve and specialized communication provided by military assets. The information management function integrates functions and data to provide the right knowledge to the right user at the right time in the right form. Battlespace awareness requires continuous battlespace surveillance with reachback to distributed resources, which in turn drives the need for real-time acquisition, fusion, and dissemination of the best available data. A global geospatial and

temporal database allowing integration, fusion, and correlation of external data must be part of an entire position, navigation, and timing system. Assurance is provided through the use of physical security and layered authentication with multilevel security. Reachback becomes a critical consideration. Because as few functions as possible are deployed to the forward base, the base must have high-bandwidth connectivity to reachback elements providing the bulk of support functions. To ensure security over all connections, end-to-end encryption is used on primary, high-bandwidth commercial systems and narrow backup links.

These enabling technologies lead to new operational concepts. The distributed Joint Forces Air Component Commander (JFACC) and air traffic control functions significantly reduce the Air Expeditionary Force footprint, accurate navigation supports all-weather close air support, and real-time dissemination enables dynamic air interdiction.

1.3.5 1998 SAB Study—*Information Management to Support the Warrior* (SAB TR-98-02): Pulling It All Together

Many common themes run through all of the reports listed in Section 1.2 *Information Management to Support the Warrior* strongly seconds many of the proposals made in the earlier studies. It goes much further than the earlier studies, though, in weaving together the threads, providing a high-level description for system implementation, assessing current technology, broadening applicability beyond command center users, and entailing more than simply networking together existing systems in a so-called network-centric fashion. This section briefly summarizes the key findings of the 1998 study to set the stage for the 1999 study.

1.3.5.1 1998 Study Findings

The 1998 study developed eleven key findings relative to the JBI:

1. Combat information requires management
2. A staff function is required to operate and manage the system
3. Human control with rule-based information decisions is required to achieve decision cycles based on information superiority
4. Data need to be organized by referencing and cataloging
5. Data need to be assembled into useful information
6. Objects can be published for common sharing
7. Subscription or search meets user needs
8. The JBI creates a common operating picture
9. Selected information requires constant updating
10. Information must be presented at the user's desired level of knowledge
11. Information validity is achieved through control of inputs

Other important recommendations of the study include:

- Continue the evolution to network-centric warfare on the way to information-centric warfare
- Leverage and explore current programs and investments
- Leverage new joint coalition operational concepts

- Build on Air Force SAB and Defense Science Board recommendations
- Present information in a use-driven fashion (the right data when and where needed)

In light of these recommendations, the 1998 study presents a high-level description of the JBI architecture. The study divides the system into three broad functional categories as shown in Figure 1-1: input, manipulation, and interaction. Information must get into the JBI, information must be manipulated to produce knowledge, and people or agents must be able to interact with the knowledge-rich results of the manipulation. The following sections describe these functions in detail.




Figure 1. The Battlespace Infosphere functions described in the 1998 study: Input, Manipulate, and Interact

1.3.5.2 Input

For information to be useful, it must be available to those who need it. Information enters the JBI from a variety of sources. While not physically present in a single system, information is present as an *object* in an *information space* within the JBI. Some specific sources for information contained in the JBI include combat support products (for example, fuels, munitions, supply, medical, and personnel systems), raw imagery as well as fused data and analysis associated with the imagery, background material (for example, maps), planning or execution products, command guidance, and user information products and databases.

Identification and authentication are key to the ability to ensure that the support products and user-generated products come from sources that are reliable and properly labeled to indicate information pedigree. Similarly, identification and authentication are used to verify that commands and execution activities are entered by legitimate users and that interactions in the JBI are conducted by people with the appropriate authority.

Access and translation technologies provide the capability to interface with existing (legacy) planning and execution systems, to translate information collected from the many input sources and transform it into a form suitable for manipulation throughout the JBI, and to support the query processes necessary for fuselets and users to select information of interest.

Categorization enables information to be characterized relative to other information.

Domain-specific taxonomies and ontologies establish a vernacular for describing relationships between different types of information. Relevance measures indicate the relative match between information content and the intent or function to which the information might be applied. Expectation-driven change-detection techniques can be used to assess whether the information being gathered is leading to a consistent state of awareness and to anticipate the arrival of new corroborating information.

1.3.5.3 Manipulate

Once information is placed in the JBI, it can be manipulated to derive new information or knowledge. Information manipulation is encapsulated within five processes: (1) The *publish* process puts information into the JBI as an object. (2) Other systems and human operators receive the published information automatically through a *subscribe* process. (3) At the same time, published information can be automatically changed into a new representation or combined with other information via a transform process. (4) A user or system that doesn't subscribe to a particular information object can still access that object using a query process, similar to a web search or database access. (5) Finally, the internal operations of the JBI can be modified and tuned using control processes.

The core features of the JBI are based on the concept of manipulation to create knowledge. This concept is composed of the ideas of publishing objects in the JBI so that they can be shared with others; subscribing to objects to be made aware of the most up-to-date information available; transforming objects into new objects, representations, or aggregate objects; allowing queries to find information within the JBI; and controlling the operation of the JBI to ensure that it is correct and robust.

The *publish* and *subscribe* mechanisms are the key to the JBI: they provide the means for communication among systems and people, and they provide a record of published information that can be queried or analyzed later. But unlike book or newspaper publishing, JBI publish and subscribe transactions can operate quickly so as to form sensor-to-shooter connections and other near-real time linkages. The single publish and subscribe mechanism suffices to provide the wide range of communication and system-integration functions needed by the JBI. Four important aspects of the design are

1. *Information objects obey standard definitions.* These objects might be likened to electronic forms, where the form is rigidly structured to record, in separate named fields, all the information required to describe the object. Objects of different types require different forms. Determining the universe of JBI object types required to support a battlespace is an important part of the JBI design. The study team recommends that the Assistant Secretary of Defense for Command, Control, Communications, and Intelligence (ASD-C³I) lead the development of common information representations across all Services.

2. *Use-driven object routing is combined with object sharing via publish and subscribe to exchange information in the JBI.* When new information objects (input from some source or fused input from multiple sources) are published, the publish and subscribe mechanism notifies all subscribers to that data of the change. JBI participants subscribe to objects by specifying essential properties of the objects they seek. An important feature of the subscription mechanism is that linkages need not be known in advance: new subscriptions can be established at any time by people or JBI processes that need information.
3. *Transformation and aggregation are performed via fuselet processing.* While the publish and subscribe mechanism routes objects from their sources to their seekers, the collection of JBI processes actually performs information-processing activities such as fusion, aggregation, and filtering. The fuselet enters one or more subscriptions in order to collect the information it needs. Whenever a new object is published that matches a subscription, the fuselet process is triggered and executed. The fuselet may examine the newly matching object and determine that it is not relevant to the fusion task for which the fuselet is responsible; subscriptions provide a coarse filter on objects, but only the subscriber can examine the details of the object fields and make decisions. If, on the other hand, the fuselet determines that it should issue new results, it publishes a new object to the JBI, which in turn may trigger other fuselets. Fuselets have many uses: they can bring information into the JBI, transform sets of JBI objects into aggregated objects, or gather objects for presentation and automatic report generation. The inputs to fuselets are typically subscriptions to JBI objects, and the outputs, where needed, are typically in the form of the publication of further objects.
4. *The JBI must include management and control functions.* These tools monitor and control such aspects as JBI establishment, performance, bandwidth allocation, security, data management, network configuration, and repair or restoration.

1.3.5.4 Interact

People and systems interact with the JBI to provide the outcomes shown in the lower portion of Figure 1-1. Operations supporting these outcomes vary in their complexity and in the extent to which they are embedded within the JBI or are services provided by the JBI in conjunction with external, connected systems. One method of interacting with the JBI is through presentations geared toward the decision maker. The display may be specific to an individual or to the position and types of decisions being made; some possible avenues of interaction include natural language, 3-D visualization, and shared workspaces.

Objects published in the JBI must be formatted for compatibility with the format expected by the person or user interacting with the information. This is especially true of legacy systems that rely on the JBI to provide input and output services to and from other information systems. In addition, there is a need to format and filter information based on the relative display requirements and information requirements of the person who needs the information. For example, while the target planner and the strike pilot need to know many identical things about an assigned target, the pilot's display is not capable of displaying all the information, so the information must be filtered and formatted to provide the appropriate display (this is also an example of a decision-centric display).

As the JBI operates, unexpected relationships that indicate useful information may be identified. This phenomenon is known as task-centered discovery. Just as an increase in the number of pizza orders at the Pentagon may have a discovered relationship to the conduct of contingency

operations, other information objects may have as-yet-unknown relationships that can drive the interaction with the JBI. A JBI for a foreign government might push data about the pizza orders to a person or system even though that information was not subscribed to or queried. These relationships will depend on the tasks being performed by the person or system and on the observed or discovered relationships.

Interaction with the JBI will depend on the users of information. Rather than broadcasting large amounts of information and expecting the people at the other end to wade through 10,000 e-mail messages and thousands of pages of information to discover the pieces they need, the JBI will forward to them only the pieces they need, without an explicit subscription or query for the information. This reduces the information overload experienced by users in a “push” system, or in a “pull” system such as a Web search engine.

1.4 Candidate Technologies

Information Management to Support the Warrior includes multiple views of candidate technologies. One view classifies candidate technologies into four categories: (1) ready to use, (2) available in 5 years through commercial R&D, (3) available in 5 years through government R&D, and (4) lacking sufficient R&D—in need of programmatic support. While this view leads to an optimistic assessment—fewer than 15 percent of the candidate technologies fall into the last category—another view in the report shows a muddier picture: a database containing 207 ongoing R&D programs shows that key enabling technologies are being covered by multiple programs, 64 percent of them by five or more programs. It is clear that overlapping programs must be coordinated and, where justified, combined, to ensure efficient use of scarce resources.

While examining technologies that may be available in the future, the 1998 study recommends building a near-term JBI relying on currently available COTS and government off-the-shelf (GOTS) technologies. Spiral development enables incorporation of new technologies as they become available. During spiral development, JBI developers should influence the evolution of immature COTS technologies to meet JBI needs.

1.5 1999: This Report—Building the JBI

While the 1998 SAB study showed the feasibility of building a robust and capable JBI, the report’s technical detail was insufficient for initiating and guiding a JBI implementation. The SAB was commissioned in 1999 to further refine the technical makeup of the JBI. This report is a result of that refinement—a technical architecture description and a roadmap for building the JBI. This report provides more detail on the technical architecture of a JBI and on the technologies being developed to support this architecture, especially commercial technologies that must be exploited. Despite the linkage between studies, this report is self-contained; the 1998 report is not a prerequisite.

1.5.1 Charter

The charter for this study appears in Appendix A. Summarized, the charter is to continue to study combat information management and develop specific recommendations that support the

implementation of a JBI. The charter focuses the study's attention in 5 areas: (1) Continue to assess the commercial research in information management technology so that the advances may quickly be applied to combat information management systems through spiral development. (2) Identify approaches for creating combat information management systems and for developing rule-based information distribution processes. (3) Identify interoperability issues to determine how to support the Service-unique and coalition (to the extent possible) combat information requirements when operating in a joint and/or combined environment. (4) Investigate and document where DoD resources need to be applied to support the military-unique requirements in combat information management. (5) Develop an implementation plan, including key demonstrations, to define a spiral development process aimed at achieving an operational capability (perhaps incrementally) as soon as technologies are available.

1.5.2 Study Composition and Approach

The organization of this study followed directly from the high-level architecture described by the 1998 study. First, three panels were created, each researching one of the three major JBI subsystems, namely, *input*, *interact*, and *manipulate*. In addition, a Joint Panel was formed to provide guidance on issues such as interoperability and joint acquisition. Finally, a Commercial Panel determined the current state of the art and important trends the JBI should follow with regard to relevant developing technologies. The members of this study bring a wide breadth of expertise—from the defense industry, defense labs, nondefense commercial sector, academia, and the military—to the table. Panel members and their affiliations are listed in Appendix B.

In doing research, panels gathered information in meetings with a wide variety of sources. These meetings are enumerated in Appendix C. While each meeting was usually initiated by a single panel, representatives from three or four panels were often present. The result is that all panels developed a coherent understanding of the military's needs, the relevant technologies, and interfaces between the various JBI subsystems. The Interact Panel met with potential users of the JBI at several military installations to determine how to best satisfy their user interface requirements. The Interact Panel also visited several research labs to evaluate the state of the art in HCI. The Input, Manipulate, and Commercial Panels met with leading technology companies such as Sun Microsystems, Oracle, and Boeing to see which technologies these influential companies will be emphasizing over the next half-decade. These panels also met with smaller companies such as Vitria and Tibco, which have middleware products that may rapidly pave a road toward integrating existing stovepiped applications.

1.5.3 Study Assumptions

Two major assumptions have been made in preparing this report: First, that current commercial developments and DoD technology deployments will yield sufficient bandwidth, connectivity, computation, and storage. Second, that the JBI information assurance and protection systems will be based on broader DoD efforts, but will also include study-generated recommendations for JBI protection and adjustment of the capability to manage degradation of information.

1.5.4 The Structure of this Report

The remainder of this report documents the results of the 1999 study *Building the Joint Battlespace Infosphere*. Chapter 2 shows the JBI from the warfighter's point of view. The discussion includes operational vignettes showing the utility of the JBI in several operations. Chapter 3 includes a description of the initialization and management of a JBI, including the use of force templates for inputting information into the JBI. Chapter 4 describes developing technologies supporting JBI interaction. These technologies put the right data in the right form at the right time in the right language and the right media. Intelligent agent technologies help determine what data are needed, when, and by whom. Agents can also infer information from users' speech and actions. Agents are also an important part of the fusion engines that process the JBI's inputs. Chapter 5 describes the JBI's management of inputs and the important technologies relevant to the input processes. Chapter 6 explains the manipulation of data in the JBI. This includes publish and subscribe implementation, fuselet activation, and middleware tying legacy applications together. The roadmap for building the JBI, including a description of many technological developments in both the commercial and government sectors, is provided in Chapter 7. Also found there is a description of the JBI's spiral development process and metrics for measuring the JBI implementation's performance. Chapter 8 provides very specific technical recommendations that summarize and build on the recommendations in the earlier chapters. Finally, Chapter 9 provides a strategic management roadmap, addressing nontechnical issues that cannot be ignored for a complex system like the JBI.

(This Page Intentionally Blank)

Chapter 2: An Operator's View of the Joint Battlespace Infosphere

2.0 The Commander's Challenge in Today's Joint Operations Environment

Today's commanders are faced with an operational environment that is in many ways more complex than in the past. The United States is no longer focused on a well-defined threat in a bipolar environment. Military planners face an almost bewildering number of threats and operating scenarios in a multipolar security environment. In addition, there are a number of other factors that increase the operational challenges. U.S. forces are more likely to have to rapidly deploy an expeditionary force to an undeveloped operational location than to employ forces from well-established overseas bases. U.S. forces must also be prepared for a wider variety of missions—from humanitarian operations, to peace enforcement, to large-scale hostilities—and they must be ready to transition quickly from one mission to another. Furthermore, expectations have been raised that combat operations will be conducted with exceptional speed, greater precision, less collateral damage, and fewer (or no) losses to U.S. servicemen and women.

To succeed in this demanding environment, a vast amount of information must be available to decision makers at every echelon of command. In fact, as articulated in *Joint Vision 2010*, the goal is information dominance. But information dominance implies more than just obtaining information; it means converting that information into a complete understanding of the situation and sharing that understanding with decision makers at every echelon at the right time, in the right format, and at the right level of detail. It also means that information systems must be capable of serving commanders throughout the spectrum of conflict—from the early days of crisis development, through the deployment and employment of forces phases, to the turnover to international civil authorities and the redeployment stage. The JBI concept was developed to achieve these goals.

To illustrate the JBI's intended scope, a few examples of its uses are provided:

- *Crisis response.* Intense efforts will be made to gather information leading to an understanding of the social, economic, religious, and political underpinnings of a conflict during its early stages of development. Much of this information will be available from nongovernmental sources. It is particularly important that National Command Authority (NCA)-level leadership, which must decide on strategic objectives and overall courses of action, and the officers in the field responsible for executing them have a common perception of the crisis environment. The JBI enables the assembly and exploitation of common databases pertaining to the nature of the conflict, which will serve as the basis for more effective vertical collaboration on courses of action between field commanders and the NCA.
- *Planning.* Assembling and thoroughly examining all potential courses of action is an important activity. If this is done well, it will build a solid foundation for success throughout the campaign. In considering alternative courses of action, decision makers should focus on the desired effects of those actions rather than the actions themselves. Good planning is supported by common databases and involves extensive horizontal collaboration among component commanders. Campaign-level models and simulations resident in the JBI will add tremendous benefit when alternative courses of action are examined in a collaborative environment.

- *Deployment.* Deploying joint forces, possibly in an environment in which the deployment itself could be opposed, is an enormous undertaking. Unit readiness data, lift requirements, basing arrangements, initial force employment options, and a thousand other details must be examined and coordinated. Tailoring joint forces for the mission is an effective concept, but it means units that have not trained together and may not have even met each other must quickly form effective teams. All this must be done under the pressures of time, intense media scrutiny, and differing Service and coalition doctrinal issues. The JBI will enable collaborative planning on the fly on a scale never achieved in past operations.
- *Execution.* Much has been said about information dominance in combat operations. Netted sensors, integrated ISR, responsive targeting and battle damage assessment (BDA), coordinated strikes, operations security, and the underlying common operational picture are necessary enablers of dominant operations. In spite of tremendous advances in sensor and communications technologies, however, too many systems are noninteroperable because they were developed independently in languages and protocols that are incompatible. The JBI concept includes both near-term and far-term solutions to this problem.
- *Coalition and nonmilitary interfaces.* In the 21st century, military forces will frequently have to work side by side with diplomatic, humanitarian, and civil authorities. Indeed, cooperation with nongovernmental forces, private volunteer organizations, and United Nations agencies has become the norm. Nevertheless, effective exchange of information with these entities is in its infancy. The enterprise integration technology that forms a significant part of the JBI platform (discussed in Chapter 6) enables streamlined collaboration with nonmilitary organizations by transforming information as needed between disparate systems.

2.1 The Importance of Information

There is little dispute about the importance of information to military forces in the 21st century. Information has always been of vital importance to military commanders. A study of military history is, in part, a study of how commanders obtained information and how they acted on the imperfect information they were able to obtain. The technologies available today, however, have fundamentally changed the way information can be harnessed to support operations. Information superiority is not simply a worthwhile goal for future commanders; it will become an essential element of mission success. It will enable commanders who master its power to make significantly better decisions in every phase of crisis and conflict.

Achieving this dominance, however, cannot be taken for granted, even though many of the individual technologies are largely at hand. The amount of information available to commanders today has increased dramatically—in both quantity and quality. But possession of large amounts of information does not necessarily result in better situational awareness or better decisions. Even with today's technology, important obstacles stand in the way of achieving information superiority. They include

- *Information overload.* Experiences from recent exercises and combat operations have shown that too often the amount of information flowing in has overwhelmed commanders and their staffs. Without powerful systems and firm disciplines in place, too much information can result in gridlock rather than dominance.

- *Lack of interoperability.* Information comes from many sources and is conveyed through many systems. Imagery systems, radar systems, logistical databases, personnel records, and economic studies are just a few of the disparate kinds of information pertinent to military operations. Understandably, most of the systems that sense and process that kind of information were developed in individual programs (or stovepipes) with their own languages, data formats, and protocols. In many cases they are not compatible with one another.
- *Immaturity in fusion.* Even when information systems are completely interoperable, it is difficult to fuse and correlate information. Pulling data together from volumes of global reporting and turning them into an accurate picture of an enemy's position and intentions is a process that can obviously be aided by intelligent systems. It is also challenging; the potential is only now beginning to develop.
- *Limits in display technology.* Humans will always make the critical decisions in warfare. Since those decisions will be based on ever increasing amounts of information, however, more efficient ways of assembling and presenting that information in highly intuitive ways must be found.
- *Legacy tactics, techniques, and procedures (TTP).* In the commercial world, businesses that invest heavily in information technology but do not change their business processes do not realize an increase in productivity. Simply adding more data to a cumbersome process is not an improvement. On the other hand, companies that change their business processes to take advantage of information technology (usually through a spiral development process) get tremendous return on their investments. In the military, business processes are called doctrine and TTP. If military forces are to harness the full power of information, new doctrine and new TTP must be found.

The JBI concept started with a recognition of the current obstacles to information dominance. It provides solutions through its ability to treat the information process as a whole.

2.2 The Joint Battlespace Infosphere

The JBI concept harnesses the power of information with a view toward making it a decisive advantage for commanders at every level. It is the means through which the concepts articulated in *Joint Vision 2010* will be implemented. It enables getting the *right* information to the *right* user at the *right* time in the *right* format in the *right* language and at the *right* level of detail.

2.2.1 Definition

The JBI is a system of information systems. It weaves together the many individual information systems that currently support operational forces. It provides an architecture for the incorporation of future data capture technologies that exploit better sensors, databases, fusion engines, automated analysis tools, collaborative planning aides, and distribution controls. It is also a disciplined process that guides the activities of the people responsible for obtaining, verifying, fusing, presenting, analyzing, and controlling the information necessary for success in any operation.

2.2.2. The JBI From an Operational Viewpoint

The JBI can initially appear complex. It is, however, both logical and well suited to the operational environment. The following descriptions are intended to help bound and define the concept.

2.2.2.1 General

- It is a system of systems.
- It weaves together the individual C² and information systems assigned to the commander who activates it. For example, if a JBI were created for an operation by a JFACC, it would integrate and coordinate the command, control, communications, computer, intelligence, surveillance, and reconnaissance (C⁴ISR) architecture analysis/planning system, theater battle management core systems (TBMCS), the Global Command and Control System (GCCS), the Global Combat Support System (GCSS), the Joint Tactical Information Distribution System, local ISR systems, and C² systems from assigned air, land, sea, and coalition forces.
- It is connected to, and interoperable with, information systems from supporting organizations. For example, systems from the National Imagery and Mapping Agency, the National Security Agency, U.S. Transportation Command, and the Defense Intelligence Agency—as well as international organizations and nongovernmental organizations that may be associated with the operation—would be made interoperable.

2.2.2.2 Information Collection

- The JBI is the commander's primary tool for ensuring that the information that is needed will be available.
- It subscribes to all of the pertinent information published by supporting systems.
- Where subscription is not sufficient, it contains or generates active agents that pull specific information from other networks.
- It centralizes the tasking of focused sensors, either through interoperability with existing sensor tasking systems or by containing its own tasking applications.

2.2.2.3 Information Storage and Fusion

- The JBI is connected to but does not duplicate all the databases with information pertinent to the operation. For example, it would have robust connections to the National Imagery and Mapping Agency archives, weather services databases, and logistics databases.
- In some cases, when incoming operational data need to be held but databases do not exist or are not sufficient, new databases can be created within the JBI.
- The JBI will accept fused data from existing fusion engines, whether they are within the JBI or contributing to it. For example, Constant Source, which provides graphical display of secret-level multisource electronic intelligence threat data, does a certain amount of fusion, and a number of Constant Source machines would fall within the JBI. Likewise, fusion services performed at the National Reconnaissance Office, the Defense Intelligence Agency, or other supporting agencies would also be accepted.
- In cases in which new opportunities to fuse information are not being performed elsewhere, fusion algorithms will be built into the JBI itself.
- This concept implies a hierarchy of fusion, one in which the most complete and coherent picture resides within the JBI itself.

2.2.2.4 Displays

- Displaying complex information under the pressure of combat conditions is perhaps the most important aspect of the human-machine interface.

- Every commander, from flight leader to four star, must have a display tailored to his or her specific purposes.
- The JBI concept recognizes that display technology is advancing all the time and that improved displays are a priority.
- Data from the common operational picture will be driving individual displays. This means that no matter how individual displays might be tailored, the common data ensure that a consistent (but not identical) view of the battlespace will be maintained.

2.2.2.5 Information Permissions and Controls

- Controlling the distribution of information is important to operational commanders. The JBI is built on the concept of providing all decision makers with accurate information and empowering them to act on it.
- Nevertheless, certain sensitive plans, intentions, and intelligence should not be available to everyone.
- The JBI includes networks operating at all classification levels, but the distribution controls needed go beyond classification categories (for example, not everyone on a Top Secret net should have access to deception plans).
- Multilevel security controls embedded in networks will be incorporated if and when the technology is available.

2.3 Effects-Based Operations

Effective military commanders have always considered their actions in terms of the effects they were trying to achieve. Military historians have marveled at the ability of Wellington, Grant, and others to obtain the precise effects they desired from subordinate commanders. In those days, orders were written by hand and delivered by messenger. In a few short sentences, these commanders could describe what was needed from a specific unit and why it was important. One of the unfortunate consequences of more rapid communications is that commanders occasionally fall into the habit of telling subordinate commanders what to do rather than what they want achieved.

Effects-based operations (EBO) is a method of defining activity and issuing orders to produce a specific effect consistent with the commander's objectives. By focusing on the desired effects, the joint force commander (JFC) and his subordinate commanders can apply the full spectrum of lethal and non-lethal weapons, including those available within allied and coalition arsenals. EBO enable an efficient strategy-to-task concept in which desired effects become the measures of merit for the objectives as stated in the commander's strategy. Thus, EBO tie the planning, execution, and assessment of attacks and operations to the target and the results of attacks according to the commander's objectives and guidance. The JBI capability enhances EBO.

2.4 Process Improvement

The JBI enables the warfighter to fully utilize information and related technologies to accomplish "full-spectrum dominance" as described in *Joint Vision 2010*. The JBI becomes an

active, intelligent environment providing a current, tailored operating picture to each commander and echelon across functional areas. In addition to acquiring and delivering the most current information to the user, the JBI feeds information directly into intuitive interfaces, decision support tools, and other aids to actively assist in conducting operations. Behind the scenes, the JBI supports information acquisition from nontraditional, open source and point-of-use data capture mechanisms, including ISR and intelligence sources and automated logistics tracking.

An example of how the JBI would support advanced concepts of joint operations provides a case in point. The JFC receives the NCA guidance and objectives from the CINC. The commander then uses the guidance and objectives to focus on potential courses of action, centers of gravity, forces available, and timing. The JBI has cataloged the sources of information on the adversary's force structure according to expressed and implicit needs for information. It collects the information in common databases, assembles it using data mining and other techniques, and presents it to the commander and the commander's planners in intuitive formats. The JBI also includes several decision support and simulation tools (by feeding the appropriate data into them directly) to help the JFC and subordinate commanders develop and analyze options.

Through a collaborative process, commanders arrive at conclusions and a shared understanding of the strategy and its underlying assumptions. The JBI provides a rich collaborative environment for communications and the sharing of multiple forms of information in real time so that commanders can develop common insights into each component's tailored operating picture. The JBI captures and incorporates the options, priorities, rationales, and details and saves them for replanning, if necessary, and for after-action reviews. At this point, the strategy and courses of action are expressed in terms of the effects they are intended to create, and assessment criteria for those effects are developed. A model-based analysis will assess whether the effects (when achieved) accomplish the objectives.

The component commanders (JFACC, the joint forces land component commander, and the joint forces maritime component commander) use the collaborative process as well as their own tailored operating pictures and decision support and simulation analysis tools to develop the schemes of maneuver that will carry out the desired effects. Drawing from situation and logistics data that are constantly updating the operational pictures, they develop sets of plans and integrated, synchronized execution strategies across the joint tactical units (wings, battle groups, and corps). The component commanders, supported by their analytical tools and simulations that the JBI automatically fills with the appropriate data, can develop multiple options and dynamic combat decision aids. These aids automatically suggest the next-priority target based on the evolving situation and available forces. Additional information is supplied into JBI decision support and analysis tools on current weather, terrain conditions, and force movement and status.

Component commanders then collaboratively provide the schemes of maneuver, including desired effects, priorities, and timing, to tactical units who must develop execution plans. In addition to their own tailored operating pictures, tactical commanders have access to detailed logistics information and resource capabilities. Communications with component commanders can be carried out in real-time or asynchronously (enabling non-emergency collaboration without requiring participants to be online at the same time) with concise information on consequences

and options that may modify the scheme of maneuver or how it is carried out. After this process is complete, the JTF commander can be confident that the right forces are being applied at the right time. Additionally, as soon as targets, tasking, and timing are determined, automatic ISR collection templates for battle damage assessment can be loaded into the JBI. These templates enable the JBI to optimally schedule scarce resources or exploit data that are already scheduled to be collected.

Beneath the JBI interface, automatic processes are developed that interpret information needs, locate information, integrate and fuse it into useful forms, and spot patterns within data, which may indicate critical occurrences.

2.5 Operational Scenarios and Vignettes

To illustrate the impact of the JBI on the military *process*—the conduct of war—this section provides an overall operational scenario, and several vignettes within the scenario, to show how the JBI enables full-spectrum dominance as outlined in *Joint Vision 2010*. This depiction of a JBI in operation is intended only to provide examples of how the power of information can change warfighting; it does not attempt to portray every JBI component that would be present in a real system.

2.5.1 Scenario

It is November 2008. A crisis has developed in the CINC, U.S. Central Command (CINCCENT) area of responsibility (AOR). The Central Intelligence Agency has issued a notice directing the intelligence community to increase emphasis on collecting and analyzing information in the Iran-Azerbaijan region. Currently available information indicates that tensions in the region are likely to lead to an outbreak of regional violence, and the United States will likely deploy forces to help resolve the situation.

No U.S. air or land forces are currently deployed to the area. Naval forces are approximately one day's sail away from the Persian Gulf. Countries allied with the United States, such as Turkey, Saudi Arabia, and Egypt, may agree to host U.S. forces. The NCA directs CINCCENT to prepare for action as appropriate. The NCA also advises CINCCENT that Russia states it will refrain from military operations as long as forces do not cross the border into Russia. A quickly eroding situation, the distance of U.S. forces from the region, and the geographic separation of units involved in the planning compress the timeline for action. As the crisis quickly escalates toward conflict, it is clear that a substantial military response by U.S. forces will be required.

CINCCENT, the JFC, initializes a JBI using predefined force posture data to speed response time to the crisis. The JBI provides a joint information structure that is tailored to the situation and fully integrates designated units. The initialized JBI is preloaded with information on available support services (movement, logistics, beddown, communications, ISR, etc.), so these services can be managed and provided to forces in the theater.

Within the JBI, information from traditional and nontraditional sources is collected and stored in the databases. The resulting surge in data is distilled into an understanding of the situation through the use of greatly improved data mining and data fusion tools. Units are added to the JBI

as they are assigned to the JTF. Publish and subscribe arrangements are established. Course of action analysis and collaborative planning sessions are enhanced through extensive use of video teleconferencing and modeling and simulation tools. Information systems used by coalition partners—some their own and some provided by the United States—are fully interoperable, and access to releasable information is managed with the JBI.

2.5.2 NCA Guidance and Objectives

The NCA provides the CINC with guidance that directs the JFC to prepare courses of action to attain the following objectives:

- Deter further action by the nation initiating hostilities and prevent the situation from deteriorating
- Employ appropriate military action to eliminate an enemy special operations force threat in and around one of the major cities
- In the event of cross-border operations by any adversary, defeat intruding forces to prevent any further combat operations
- Conduct offensive operations as appropriate to neutralize a mobile theater ballistic missile (TBM) threat thought to be capable of delivering weapons of mass destruction (WMD)
- Restore regional stability

2.5.3 JFC Intent and Rules of Engagement

When the JBI is initialized, the JFCs provide their component and subordinate commanders with their intentions and rules of engagement (ROE) so that detailed plans to attain NCA objectives can be developed. These inputs can be modified as necessary as developments warrant.

- Component and subordinate commanders coordinate a plan to
 - Gain air superiority
 - Direct the 82nd Airborne Division to establish an airhead near a major city
 - Conduct continuous surveillance of theater from airborne sensors
 - Interdict mobile TBMs
 - “Fix” forces in the south, denying their support to operations in the north
 - Prepare and execute a prehostilities information warfare campaign
- Should deterrence fail, commanders halt invading forces using the ROE, provided
 - Forces in or flying over designated areas will be attacked
 - Land forces are free from attack as long as they stay within their borders
 - Forces supporting designated hostile activities will be attacked
 - Integrated air defense systems will be targeted and attacked
 - TBMs will be targeted and those tactically deployed or in flight will be destroyed

2.5.4 Vignette 1: Deception Operations to Deter Hostilities

Reports from a variety of intelligence sources—fused and coordinated within the JBI—provide JFCs with a clear picture of what information each of the belligerents is using to piece together

U.S. activities and intentions. For the most part, the adversaries are building fairly accurate pictures based on a variety of collection methods. U.S. intelligence sources and methods are sensitive and must be restricted to the highest levels of command. The joint commanders and the Joint Staff, however, are presented an opportunity.

An information warfare plan that depends largely on ensuring that the adversaries correctly receive signals about the U.S. resolve to act—supported by a focused deception plan to keep them off balance regarding specific force deployments and intentions—is created to discourage the initiation of hostilities. The JBI enables the coordination between the intelligence community and commanders in the field necessary to execute a plan of this kind. Access to and dissemination of sensitive information is controlled within JBI protocols.

2.5.5 Vignette 2: Insertion of U.S. Forces to Establish an Airhead (Deter)

Before deployment of the planned U.S. deterrent force can be fully executed, tensions in the region increase. Border skirmishes and engagements between factions in urban areas escalate. Instability in one of the countries has prompted its government to request that U.S. forces seize and hold a key airport that is necessary for humanitarian support and possible military deployments.

A task force has been formed around the 82nd Airborne Division for this purpose. The environment is potentially nonpermissive. Some tactical air forces have already deployed, but air superiority has not yet been fully achieved. Speed, surprise, and effective coordination must be attained despite the hasty circumstances in which the operation must be planned.

The existence of the JBI, together with the associated changes in C² concepts, significantly reduces the risks inherent in an operation of this kind. Both the decision authorities and the executing units have much better situational awareness. The communications gap that has existed between them for so long will be effectively bridged. This new level of situational knowledge can reduce the need for planners to include extensive hedges against uncertainty. When the executing forces on scene have the kind of tactical awareness that enables informed decisions on the fly, less-detailed orders are needed. The commander's intentions are fully understood, and combat actions become self-synchronizing.

In taking down this airfield, the JBI provides the standardization and common reference point for maps, terrain data, geospatial reference systems, threat data, and the host of other information needed to execute the mission. It also allows the synchronization of temporary air superiority, early warning denial, execution of deception plans, search and rescue operations, and other elements of the mission.

As this vignette makes particularly clear, the JBI must be joint. Full exploitation of the information concepts established in *Joint Vision 2010* and enabled by the JBI can be realized only when joint doctrine, organization, training, materiel, leadership, and people are adapted to the new environment.

2.5.6 Vignette 3: Elimination of the TBM Threat

As the deployment of U.S. forces gets under way, one of the belligerents threatens use of its TBM force to halt the operation. The country possesses a significant number of missiles plus a small number of nuclear, chemical, or biological warheads that are believed to be operational. This WMD threat is having a substantial impact on the political and media environment in which strategic decisions must be made.

A theater missile defense plan has been developed. Besides the traditional pillars of theater missile defense, the plan includes strong political statements resolving to hold any user of WMD strictly accountable. The JFACC set a goal of doubling the effectiveness of attack operations and has put together a plan for the conduct of attack operations that takes full advantage of the JBI. Extensive intelligence preparation of the battlefield has been conducted. Every scrap of information that could contribute to indications and warning of impending TBM operations has been examined. Human intelligence, signals intelligence, measurement and signature intelligence, and imagery intelligence indicators have all been analyzed; potential hide sites for transporter-erector-launchers have been mapped using terrain delimitation and other techniques. Potential ground movement patterns have been programmed into Joint Surveillance, Target, and Attack Radar System (JSTARS) and Discoverer II computers. Prescribed responses have been stored in the JBI. A network team devoted to the problem has been identified, and publish and subscription protocols have been tailored to ensure the most rapid and complete sharing of information.

Suspicious ground movement is detected by one of the on-station JSTARS. Following the fuselets created for this situation, a number of activities are simultaneously put in motion. Imagery and electronic intelligence sensors are immediately alerted, retasked, and netted. Two nearby F/A-18s and two EA-6Bs are diverted from a routine strike mission. Should attack operations prove unsuccessful, winds and weather effects are calculated. Based on bits of corroborating data from Rivet Joint and Predator, an engagement decision is made. Several sensors detect surface-to-air missile activity, and the EA-6Bs provide protective cover. The F/A-18 strike is successful and a TEL is destroyed before launch operations can commence. The netted sensors that have been temporarily tasked to support this contact continue surveillance until sufficient BDA is collected to confirm the kill. Enemy order of battle is adjusted accordingly.

Particular note should be taken of the C^2 concepts behind an operation of this nature. Many decisions are required before the action begins. Publish and subscribe protocols must be put in place to ensure that information is available where it is needed, and fuselets must be created in order to prescript coordination to the greatest possible degree. ROE must be published, as well as rules that govern the preemption of forces from other missions. These are all procedures that enable virtual teams to be formed on the spot as circumstances dictate. During execution, however, decisions are delegated to tactical levels. In this case, any one of a number of on-scene units could have made the decision to engage—a decision based not on a visual identification, but on the aggregated weight of disparate pieces of information. This further points out that the JBI concept rests on more than just information systems; it depends equally on advanced training and doctrine.

2.5.7 The JBI as Seen Through the Operational Scenario and Vignettes

The scenario and vignettes highlight how the JBI, for the first time, can fully exploit information dominance. First, the JBI allows better preparation and collaboration among a variety of entities and sources across geographically disparate locations. Second, the JBI enables better synchronization of battlefield activities based on the shared understanding enabled by information from traditional and nontraditional sources. Furthermore, the JBI permits higher-level, accelerated analysis of alternative courses of action to occur in time-compressed situations.

2.6 The JBI Enables New Ways of Conducting Operations

As the scenario and vignettes should make clear, a commander's information system must serve the commander's process. Too often in the past, operational processes were constrained by system limitations. The flexibility and pervasiveness of the JBI, on the other hand, offers the possibility of entirely new ways of doing business. In fact, the power of the JBI cannot be realized without improved processes. In the business community, experience has clearly shown that upgrades to information systems without a parallel business process improvement are not successful. As the JBI is integrated into operations, new ways of conducting operational business that can't currently be envisioned will evolve on an expanding scale.

2.7 The JBI From a Joint Operational Viewpoint

The JBI is the necessary ingredient that will allow JFCs and subordinate air, land, sea, and coalition force commanders to conduct decisive operations according to the precepts set forth in *Joint Vision 2010*. As stated in *Concept for Future Joint Operations: Expanding Joint Vision 2010*, "JV 2010 is built on the premise that modern and emerging technologies—particularly information-specific advances—should make possible a new level of joint operations capability." The JBI enables information superiority by leveraging innovation in information technologies and military doctrine and TTP. As envisioned and demonstrated in the scenario and vignettes in this chapter, the JBI permits information operations that effectively achieve full-spectrum dominance over an adversary and control any situation through

- *Dominant maneuver.* The multidimensional synchronized application of information to permit joint engagement, and mobility capabilities positioning and employing widely dispersed joint air, sea, land, and space forces to accomplish the assigned operational effects.
- *Precision engagement.* An information system of systems that enables forces to locate the objective or target, provide responsive C², generate the desired effect, assess the level of success, and retain the flexibility to reengage with precision when required, coordinated to a common joint battlespace environment.
- *Full-dimensional protection.* An information-integrated, multilayered offensive and defensive capability to protect forces and facilities at all levels from adversary attacks while maintaining freedom of action during deployment, maneuver, and engagement.
- *Focused logistics.* The fusion of information with logistics and transportation technologies to provide rapid crisis response, to track and shift assets even while en route, and to deliver tailored logistics packages and sustainment at the strategic, operational, and tactical levels of operations.

2.8 Factors That Must Guide Development

Information is not an end unto itself. Information is valuable only as it aids in achieving mission objectives. As the Air Force begins to implement and refine the JBI concept, several factors must guide development of the JBI and are worthy of note:

- *Interoperability is essential.* Joint operations are the normal way U.S. forces conduct business. Moreover, coalition operations are more and more becoming the norm. In addition, the information necessary to a commander comes from a wide variety of sensors and sources. For an information system to be useful, it must seamlessly interact with all other systems—new and legacy—involved in the operation.
- *Decisions at all levels must be supported.* Important decisions are made at every echelon of the command structure, from the flight line crew chief to the JFC. The JBI must support information needs at all levels.
- *All dimensions of warfare are important.* “Amateurs talk tactics; professionals talk logistics.” This old axiom holds true today and reminds the listener that focusing on combat operations alone does not lead to success in war. Precision strike, sensor-to-shooter, and similar combat concepts have received considerable attention in recent years. But good logistical planning, intelligence preparation, and other aspects of military campaigns, while sometimes less glamorous, also depend on information and must also be served by the JBI.
- *Systems engineering is required.* Although information systems are getting more interoperable each year, a system of this complexity will not come together on its own. Good systems engineering decisions will be required through the life of the concept.

Finally, it is important to note that spiral development is the best (and only) process by which the JBI can be matured as a capability. Spiral development is a concept in which development of the JBI is tightly coupled with the development of innovative TTP (new business processes). This approach to development and acquisition is essential if the JBI is to be realized.

2.9 The Imperative to Act

Moving toward this advanced concept for harnessing information is not an option. Continuing investments in ISR means more information is collected. However, there is no point in more and better ISR if the data can't be handled and exploited.

The Internet opens access to nontraditional sources that can be invaluable to forces conducting operations. However, the Internet can also compound the problem by overwhelming warfighters with the sheer volume of 1s and 0s.

If joint military operations are to succeed in an increasingly complex, information rich world—adaptation and decisive action are needed. Developing the JBI is the first step in this direction.

2.10 Summary and Transition

Chapter 2 draws a picture of how the JBI supports warfighters and leads to decisive information superiority for U.S. forces. In the following chapters, this report moves from the operational view of the JBI emphasized in this chapter to a technical view. The first step in this transition,

which is found in the next chapter, is the description of the operation of the JBI. The subsequent chapters describe in detail the technologies needed to perform the interact, input, and manipulate functions of the JBI. Chapter 7 examines existing technologies, and Chapter 8 provides a roadmap for development of the JBI.

(This Page Intentionally Blank)

Chapter 3: Operating the JBI

3.0 Introduction

The previous chapter presented the benefits of using JBI from the warfighter's perspective. Before transitioning to a technical perspective of the JBI, this chapter describes the operation and management of the JBI throughout its lifecycle. It covers the infrastructure needed to support the JBI, the information staff needed to manage the JBI, peacetime use of the JBI, and the standing up of the JBI for a crisis or contingency. A revolutionary concept presented here is the notion of a force template, the collection of information objects a unit brings to the JBI when the unit joins the assigned JTF.

3.1 The JBI Infrastructure

The JBI is a complex information management system that integrates many existing and future systems. As shown in Figure 2, the JBI forms an umbrella under which GCCS, GCSS, ISR platforms, and other systems work. That is, the JBI does not replace the systems, but it is built on top of or around them. The JBI integrates systems, moves information between them, and ultimately provides users with the information they need. In one sense, the existing systems and those developed as the JBI comes online form an application infrastructure for the JBI.

On the physical and networking side, the Global Grid is the foundation for the JBI. Worldwide, high-bandwidth communications provide the connectivity needed by the JBI. The JBI has a small forward footprint, and it provides many of its services via reachback connections. Much of its computational power and storage capacity can reside in the rear, with a minimal number of components near the battlefield.

As one of its management functions, the JBI monitors the available communications bandwidth. If some data links become degraded or lost, the JBI reduces information flow so that no links become saturated.




Figure 2. The JBI integrates existing and future C², ISR, and support systems

3.2 Managing the JBI

While the JBI requires a physical and computational infrastructure, the human infrastructure needed to operate the JBI is absolutely essential. A skilled, well-trained staff of professionals is needed to ensure that the JBI supports the operational commander and all users in accomplishing the mission.

The information staff represents a new element in the organization structure of the operations center. The chief information officer (CIO) reports to the commander. The CIO is responsible for establishing and maintaining the JBI and is supported by an information staff. The CIO advises the commander on the capabilities of the JBI and the ways information can be brought to bear in the current contingency. The CIO ensures, through the information staff, that the commander's information needs are met. The CIO must control JBI reconfigurations made by the information staff on behalf of other users to ensure that JBI integrity is maintained. The CIO must be both warfighter and technologist, understanding the information technology behind the JBI and knowing how to apply that technology in battle.




Figure 3. The information management staff organization. The staff includes functional expertise from ops, intel, logistics, and communications

3.2.1 Information Staff Functions

The functions of the information staff include

- Information management
- Information protection
- Information infrastructure
- Information operations

These functions are described in the following sections. Note that the division of responsibility listed here is arbitrary. As the JBI is deployed, the proper information staff structure will evolve to fit the business rules of the JBI.

3.2.1.1 Information Management

The information management staff function is responsible for maintaining the top-level JBI architectural structure. In the implementation of a crisis-specific JBI, the information management staff assures that appropriate data populate the JBI and adhere to an agreed ontology. The librarian function reports to the information manager. The information manager has responsibility for data quality, JBI health monitoring, and coordinating the implementation of “business rules” within the JBI. The information manager is the approval authority for implementation of new features within the JBI.

3.2.1.2 Information Protection

The information protection staff function is responsible for information security, including security policy, attack detection, countermeasures, and access control for JBI users and client machines and programs. Security policy is set at the highest levels, and the information protection staff implements the policy through the use of low-level security mechanisms such as firewalls and access controls. This staff will deploy both host-based and network-based intrusion-detection systems to ensure that intruders do not go unnoticed in the JBI. They will investigate a variety of alarms: those set off by the JBI’s access controls when a user tries to

access unauthorized data, as well as those set off by the intrusion detection system. If an intruder is found, the hole through which he or she entered must be found and closed, and any damage done repaired. Alternatively, intruders may be diverted into a “sandbox,” in which their actions are controlled, their attacks have no impact, and false data may be planted in the hopes of deceiving the enemy.

3.2.1.3 Information Infrastructure

The information infrastructure staff function is concerned with transportation, setup, operations, and maintenance of facilities, computers, and LAN hardware. The information infrastructure staff initialize the JBI with network details so that the JBI is aware of its configuration and can perform bandwidth management. They ensure that adequate computer and storage resources are allocated to operate the JBI. This functional group performs low-level operating system administration duties such as making backups and protecting copies. They monitor performance and ensure that it meets mission needs, they update software as needed, they diagnose and repair system outages, and they consider physical environment issues such as electrical power and air conditioning.

3.2.1.4 Information Operations

The information operations staff function is responsible for managing the day-to-day operations and administration of the JBI and the tools that are used to operate the JBI. The information operations staff oversees the generation and validation of new features and provides expertise in configuration and operation of the software tools—the foundation of JBI functionality.

3.2.2 Information Staff Technical Expertise

JBI information staff members have expertise in a number of technical areas, including

- *Web operation.* Potentially, much of the JBI platform will be built on web technology.
- *Information design.* Objects will be used to represent information, requiring an understanding of object technology and information webs (such as ontology, schema, and structured common representation [SCR]).
- *JBI scripting language.* Some programming expertise will be needed for fuselets, agents, and other forms of JBI manipulation.
- *C² business logic.* This knowledge includes military business rules, command guidance, and mission plans.
- *Information assurance.* Expertise in the most current defensive techniques will be essential to ensure availability of JBI functions.
- *User modeling.* Data must be formatted and filtered to meet the needs of warfighters.

3.3 JBI Lifecycle

Given the people and systems that make up the JBI, this section describes the JBI lifecycle. That is, how a JBI is stood up, how it grows, and how it is maintained.

3.3.1 Peacetime Operations

A JBI is established at the direction of a JFC to address a particular crisis or contingency. During peacetime, one or more JBIs will be in operation for testing and training.

3.3.1.1 Units in Standalone Operation

During peacetime operations, each unit operates its own information systems, which comprise a *Unit Infosphere*, for day-to-day operations. These systems are compatible with a JBI. They use the same information object standards and information structures of a JBI, but they operate in isolation. For example, a wing's Unit Infosphere may implement many of the information flows required within the wing's structure.

3.3.1.2 JBI System Testing and Development

The JBI will continually evolve as new technologies, features, and operating procedures are introduced during its spiral development. Updates to a complex system like the JBI cannot be tested in a vacuum. A full JBI must be established to ensure that updates perform correctly without adversely affecting the existing functionality.

The development of operational concepts must go hand in hand with the spiral development of the JBI. As new JBI software and systems are tested, operating procedures must be updated and validated. Exercises such as the Joint Expeditionary Force Experiment (JEFX), in which actual operators evaluate the systems and procedures, ensure that the system meets the users' needs before a new JBI version is deployed.

3.3.1.3 Training and Exercises

Once a new version (including both information systems and operations procedures) of the JBI is deployed, all users must gain experience working with the JBI. Some training is performed locally using the systems comprising the Unit Infosphere. Large-scale exercises are facilitated through the standing up of an exercise JBI. Because the JBI does not require that all units be physically present in one location, geographically dispersed units can plug their Unit Infospheres into the exercise JBI from their home bases using force templates, described below. All users, no matter what their functional areas, get to use the JBI as they would in an actual crisis. In addition, their inputs on JBI operations can be used to improve the system or processes that make the JBI work.

3.3.2 Contingency and Crisis Operations

While the JBI supports training and exercises, its real purpose is for a contingency or crisis. This section describes the operation of the JBI in such a situation.

3.3.2.1 Standing Up the JBI

When the CINC decides that a situation in the AOR has escalated to the point that a JBI is needed, the JFC's information management professionals stand up a JBI that is focused on that crisis or conflict. When initialized, the JBI contains numerous default settings, processes, and data values, so all users work with a common operating picture right from the start. Many of the

needed data values can be predicted before a crisis because the AOR is well understood before the JBI is stood up. Some data values that may be part of the initialized JBI are

- *The JFC's objectives, ROE, and operations concepts.* This information provides the JBI with the objectives it is supporting and the constraints within which it must work.
- *Regional and background information.* The JBI contains maps for the region of concern. Weather data—both long-term climate information and short-term forecasts—are critical to JBI users and will be linked into the JBI. Cultural characteristics of the region may provide critical information. The adversary's culture may determine the effects the JFC wants to impose to favorably end the crisis. Understanding of the local allies may foster better coalition operations.
- *Adversary information.* In addition to the cultural information mentioned above, the JBI is preloaded with potential targets and centers of gravity. These are developed through exercises and precrisis planning. The adversary's political structure is another key piece of information that must be available to the JFC.
- *Assigned unit information.* When the JBI is activated, the JTF has a variety of units assigned to it according to their location or functionality. The JBI is preloaded with these units' capabilities and logistics requirements. In addition, the units' information requirements are loaded into the JBI as subscriptions. Any ISR information the units can provide is also loaded into the JBI.
- *Available ISR resources.* At startup, the JBI is linked to worldwide imagery databases. In addition, it is informed about sensors that can provide near-real time data and any requirements for tasking these sensors. As described in the preceding paragraph, many units bring their own ISR capability, and the JBI tracks that as well.
- *JBI server configuration.* Standing up a new JBI requires identifying and configuring a set of servers to provide JBI platform services, including publish, subscribe, and query. Missions of significant size will probably require several sets of servers, and a mission may have several different sets of "rear" and "forward" servers. These servers will operate together, as a "federation," to serve as a single JBI.
- *Information access privileges.* The JBI must be able to give a security classification to information on the basis of the sources through which data are input, fuselets that execute to generate the information, or explicit labels provided by authorized users. It must also know which users, machines, and networks have permission to access different levels of sensitive data. A newly initialized JBI will be configured with a default classification and privilege scheme, and the information staff will update this scheme to fit the commander's needs.
- *Publish and subscribe requirements.* These provide the means for communication among users and systems from assigned units.

The default processes would include both heavyweight processes, such as TBMCS, and preloaded fuselets. Fuselets enter one or more subscriptions to collect the information they need. Whenever a new object is published that matches a subscription, associated fuselets are triggered and executed. Fuselet execution normally generates new information, sometimes by triggering chains of fuselets. The default JBI fuselets would subscribe to default JBI information.

Once the JBI is established, clients connect to the new JBI by simply naming it (similar to using a Website name). Software clients of the JBI can locate all the necessary JBI platform services

using this name. Users will access the JBI using a “JBI browser,” analogous to a Web browser, which can be used to view common information objects and to issue queries to the JBI.

Users will need to customize their JBI configuration. For example, the JFC may require frequent updates on the status of a key enemy target, so the information staff will generate a new fuselet subscribing to relevant data. Similarly, information access privileges may need updating to reflect the needs of the operation. High-level directives such as the ROE may also be updated throughout the JBI’s existence.

While users and the information management staff may manually tweak the JBI, major reconfigurations are too complex to be done manually. As new units, many of them with their own information systems, join the JTF, the JBI must instantly reconfigure itself to incorporate the unit, namely its existing computer and communications systems, users, weapons, and ISR assets. The mechanism through which units attach themselves to the JBI is the *force template*, and it is described in the next section.

3.3.2.2 Growing the JBI—Templates Let Units “Plug In”

When a new military unit is added to the mission, its C² systems must be integrated into an already-running JBI. For example, if a new unit of an advanced tactical missile system joins an ongoing mission, the systems of the unit must be linked into the JBI. From another perspective, the information structure that represents unit operations must be integrated into the information structure of the JBI.

This integration does not bring the Unit Infosphere completely under the control of the JBI. Rather, the integration is an information interface or a definition of data to be exchanged. This integration must occur quickly, with minimal human intervention. Furthermore, it must be a complete “handshake” so that the JBI knows exactly what services and information the unit provides and what information the unit needs (its subscriptions). The mechanism employed by the JBI for integrating a new unit is the force template.

The force template is a software description of a military unit that may be integrated into a JTF. Included in the force template is information about the Unit Infosphere that is to become part of the JBI. Several varieties of force templates are used. One form of force template describes a fighting unit (for example, a tank battalion or fighter squadron). The key elements of information included in such a force template include

- Force employment capability
- Ammunition inventory
- Fuel requirements
- Communications requirements
- Computing systems
- Information requirements (the subscriptions of the unit’s clients)
- Information products (objects to publish during and after the instantiation [for example, ISR])
- Personnel requirements

Another variety of force template describes a support unit. A template for these units always includes the following information:

- Information requirements (the subscriptions of the unit's clients)
- Information products (objects to publish)
- Communications requirements
- Computing systems

However, the specific publish and subscribe exchange described in the force template varies with the type of unit. For example, an airlift control unit force template would publish information about the locations and movement of relevant airlift assets while subscribing to airlift needs generated by the JBI.

The information management staff for a unit deploying to the JTF ensures that its force template is current. When the unit is assigned to the JTF, the force template is input to the JBI. Note, though, that the handshake that brings a new unit into the JBI must go two ways. The unit provides a detailed description of itself to the JBI, but the Unit Infosphere must be reconfigured to work under the JBI. Therefore, the JBI may need to provide a template to the Unit Infosphere. This template includes infrastructure information (for example, routing in the JBI networks), information that the Unit Infosphere must publish, and subscriptions that the JBI mandates for the unit.

3.3.3 Maintenance

The information staff is responsible for correct operation and maintenance of the JBI. The staff's routine duties include the following:

- *Nonstop operation.* The information staff will need to devise configuration and operation policies to ensure the nonstop operation of the JBI, a mission-critical system. For example, most data will need to be stored redundantly, with at least one protected copy. The staff will need to monitor JBI operation. When outages occur, due to equipment failure or battle damage, the staff must diagnose and repair the system.
- *Integrity guarantees.* The information staff will conduct audits and other routine procedures to ensure that the JBI platform and its clients are functioning correctly, that no unauthorized clients or users are being served, and that the information flows of the JBI are supporting the mission. Breaches or corrupted information objects must be repaired.
- *Provisioning.* The information staff must ensure that adequate computer, storage, and communications resources are allocated to operate the JBI, configuring new resources as necessary. The staff may need to adjust configurations to ensure that the JBI's performance meets mission needs. For example, each publication of an object is assigned a priority, which the staff can adjust to ensure that critical objects are processed first.
- *Supporting clients.* It is likely that the information staff will support major JBIs, such as fusion engines and C² suites (for example, TBMCS and GCCS). Support will include administering computers and networks and monitoring continued operation.
- *User admission.* JBI users must be known to the JBI in order to enforce security and access controls. The information staff will need to register new users and remove those no longer working

with the JBI. It is not essential that this function be handled by the information staff; an alternative is to delegate it to commanders or security staff through the chain of command.

- *Access control.* The JBI must be configured to enforce the access controls required by the mission's commander. For example, the CINC, in the case of a theater JBI, may allow certain kinds of information flow within the theater but restrict access from outside. Access controls pertain to all clients—users as well as clients operating autonomously (for example, fusion engines). The information staff must also defend against security breaches or attacks on the system.
- *Software updates.* The information staff is responsible for upgrading JBI software if necessary, without disrupting JBI operation. It is simpler if a JBI can be deployed with software of uniform vintage: servers are initially loaded with consistent software, and clients are allowed access only if they are sufficiently up to date. However, because some JBIs need to function over extended periods, the information staff may need to upgrade software during operation.

3.3.4 Adapting the JBI to the Mission

If a commander requires information to be collected, fused, or presented in a new way, information staff specialists will design and implement the changes to the JBI. These adaptations may require introducing new information object types, changing the configuration of existing JBI clients, or creating and deploying new fuselets. In some cases, new technology may become available during the course of a mission, and the information staff must integrate this technology into the JBI information structure. The information staff must be familiar with the information structure and its expression in publish and subscribe linkages in order to adapt the system.

3.3.5 Sharing With Coalition Partners

Interoperability is a significant challenge for joint operations. For coalition operations, interoperability appears to be an insurmountable problem. The JBI's use of force templates can significantly improve this situation. The format for a force template can be shared with joint and coalition partners so that all participants in the AOR can interface, at least to a minimal degree, with the JBI. When coalition members bring their forces to bear in the conflict, the JBI's information management staff can input force templates on the members' behalf. Some allies may already bring force templates with them. Indeed, some may have systems already configured to exchange information with the JBI. Partners with noninteroperable systems cannot immediately publish or subscribe, but their capabilities and other important information will be available to the JBI through manual entry. As the conflict continues, the information staff may install and configure a suitable "gateway" that provides access to the JBI, subject to constraints ordered by the commander. A gateway is bidirectional: it subscribes to relevant JBI objects, converts the information into the partner's format and language, and sends it to the appropriate command, control, and communications (C³) system of the partner; it also extracts information from partner systems, converts it into JBI information objects, and publishes them in the JBI. Gateways will usually require extensive configuration to describe the rules that govern information transfers.

3.3.6 Units Relieved From the Assigned Mission

When a unit is reassigned from the JTF mission, the JBI must be adjusted to operate correctly in the absence of the information objects that are being supplied by the departing unit. Essentially,

the JBI is told to “unplug” the force template provided by the unit. The Unit Infosphere for the reassigned unit is returned to its peacetime configuration.

3.3.7 Standing Down

When a mission is complete, its JBI can be largely stood down and redeployed. However, the repositories of objects collected during the JBI’s operation will usually be saved for archival value. The JBI itself may continue to operate in a severely curtailed way: repositories will still operate and respond to queries.

3.4 Summary

This chapter describes the operation of the JBI from the information operator’s view. It describes the infrastructure needed to make the JBI work. It introduces the information management staff and describes the variety of functions they perform, from top-level management of the JBI to information protection. The JBI lifecycle—from standing up the JBI to dismantling it after a crisis—is also described. The key concept introduced in this chapter is the force template, the means by which a unit’s information systems are introduced to the JBI when the unit joins the associated JTF. The next three chapters move from the operator’s view to the implementer’s view. That is, they describe the technical details for the interact, manipulate, and input segments of the JBI.

Chapter 4: Warfighters Using the JBI

4.0 Introduction

The vignettes in Chapter 2 highlight the ways the JBI supports the user with task-specific displays. Chapter 4 goes into much more detail, showing how the JBI provides the user with the means for making decisions by formatting data for individual tasks. As shown in Figure 4, the JBI provides information in a use-driven fashion (that is, when the information is needed) and even helps the user to discover information not explicitly stored in the JBI. This chapter doesn't describe JBI functions only from the user's point of view but also from a technological point of view. The chapter closes with recommendations specific to the interact segment of the JBI.




Figure 4. Functions of the Interact Portion of the JBI

4.1 Interaction Overview

The primary goal of the JBI is to provide the right information, to the right person, at the right time, using the right media, language, and level of detail during the planning, command, execution, and combat support of missions. How the JBI meets this goal is described in Section 4.2. At each step in the example, the user, environment, and task are described. The interaction of the warfighter with the JBI is then described followed by a discussion of the three categories of interface technologies (data capture, presentation, collaboration) applied to during that step.

The example is a small portion of the vignette presented in Chapter 2. To summarize, a transporter-erector-launcher (TEL) is detected in the commander's AOR. The commander orders

its destruction. A distributed staff of personnel develops a plan to carry out the commander's order. Combat support personnel prepare the assets needed to carry out the plan.

4.2 Command

The JBI user illustrated in this section is a commander inside a permanent building with full power and communications connections. His or her task is to enforce a treaty denying use of weapons to Serbian forces. A TEL is detected and the commander orders its destruction. The JBI supports task-centric decision making.

As shown in Figure 5, the JBI is stood up, or instantiated as part of the routine operational role of the commander to monitor the AOR and other geopolitical interest topics. The commander will "subscribe" to the JBI information and data sources that provide current status of activities associated with this AOR.




Figure 5. Commander subscribes to current status information to monitor the area of responsibility

Alerts relevant to the AOR may come in the form of messages, analysis reports, news bulletins, or other means, and on receipt, will be transformed by the JBI to meet the tasks and preferences of the commander (see Figure 6).




Figure 6. JBI transforms alert to support commander's task and preference

4.2.1 Data Capture

The human interface of the future will be characterized by more natural modes of interaction than are currently possible. Speech recognition, eye movement tracking, gesture recognition, and handwriting analysis will be key components of a naturalistic interface. The thrust of the collection of capture technologies is to reduce the burden of acquiring and maintaining data and providing user control input for system applications. The objective is to either collect the data in some automatic way, or to at least make the collection or input methods more natural and tailored to the task at hand. These technologies include speech processing, natural language processing and dialog, and multimodal interfaces such as speech and gesture.

4.2.1.1 Conversational Speech

A speech interface will provide a key mode of interaction with the JBI. There are two important reasons speech is so important. First, speech-based input provides a convenient hands-free method to interact with computer applications. Second, with proper design, speech input and data manipulation can be faster than more conventional computer interaction methods. Automated speech recognition (ASR) is the main technology embedded in a speech-based data capture system.

ASR technologies have significantly matured in the past several years. Due to the dramatic increases in processor speeds and memory availability, real-time continuous ASR technologies are becoming more common and will soon be a preferred human-computer interface technology. Computer C² services enable users to perform navigation and data entry functions.

4.2.1.2 Natural Language Processing

ASR enables natural language dialogue between a human and a machine. Natural language processing extends this technology to include a context- and task-based understanding of user intent that considers goals, preferences, and workflow. These range from simple C² systems to

“open window” systems that do not require users to speak in a set vocabulary and syntax. Dialog management systems model the user, have conversational skills, and can clarify and explain information. Translingual communication enables users with different native languages to communicate.

As previously mentioned, one of the aims of this collection of technology is to provide natural interfaces. A real problem with current C² systems is the level of training and specialized support required. Natural language processing in an interface can greatly reduce specialized training. This does not mean that the users are not trained for their tasks but assumes that they are trained in the language of the task and that it is now a natural language for them.

4.2.1.3 Multimodal Interfaces

A new generation of multimodal systems is emerging in which the user will be able to employ natural communication modalities, including voice, hand- and pen-based gestures, eye-tracking, and body-movement, in addition to the usual graphical user interface technologies. To a large extent, systems that fuse gesture and speech already exist, as well as those that fuse speech and head movement for hands-free HCI.

4.2.1.3.1 Speech and Gesture

An important class of technology that can have real impact on the user’s ability to effectively interact with the JBI is speech and gesture technology. There are three important advantages that the combination of speech and gesture provide. First, it can enable the user to employ different elements of communication in their most natural form, such as in “move the trucks to here.” Second, multiple input modes can significantly reduce the error rate for capturing commands to user applications. The gesture recognition can help clarify the speech input and vice versa. Finally, multiple modalities exploited simultaneously have been shown to increase efficiency and reduce time in user workload and flow.

Systems that fuse gesture and speech already exist and are being used to help structure visualizations, control military simulations, manipulate virtual objects, and designate targets in fighter aircraft. The JBI could create portals or interaction frames for the user to interact with these types of information within the JBI. The interaction frames are likely to exploit advanced visualization. The JBI user should be able to directly interact through these visualizations, either by pointing or gesturing. For example, if the commander wants to reposition some assets, the commander can gesture on the appropriate interaction frame.

4.2.1.3.2 Domain-Specific Gesturing

It is widely believed that as the computing, communication, and display technologies progress even further, HCI tasks will become the bottleneck for the effective capture of relevant information and its utilization. For future HCI techniques to be useful, they must exploit user context, preferences, and workflow. Domain-specific gesturing techniques specify and exploit a “gesture vocabulary and grammar” that can be used to help machines understand gestures specific to users, work tasks, and spoken words.

Much as shorthand can be an effective means for rapid communication, there is an opportunity for gesture recognition systems to be trained to gestures that are either already taught or practiced as part of a task. It is also possible to develop new classes of gestures that will be enabled by gesture recognition technology.

4.2.2 Drill Down

No matter how smart the presentation software is, there will always be the requirement to provide a capability for users to drill down for verifying source data and assessing accuracy and timeliness as part of the fusion process. Invariably, users will need to see some aspect of the detail behind the information presented, and that will likely be experience specific. The idea is to capture as much of the context of the interaction as possible and to enable the system to do the best job possible in presenting the next layers of detail.

4.2.3 Presentation

Presentation is the medium and format for the information that is input to or output from the JBI. Presentation technologies that potentially apply to the JBI environment appeal to the natural human senses. These include visual, aural, and haptic modalities as well as nontraditional techniques that augment the perception of human reality. While the data content of presentations will always be the most important consideration for the user, the methods and hardware devices that limit or augment the conveyance of the underlying information are increasingly important especially considering the drive toward user mobility in future warfighting scenarios. This trend toward mobility will increasingly drive the need for small, lightweight, and even wearable hands-free information presentation devices.

4.2.3.1 Personal Display Devices

Ideally, individuals can interact with the JBI data environment in free space in an unencumbered fashion. Several technologies may have a significant role in JBI interactions over the next few years. The following discussion presumes the simultaneous existence, availability, and use, within the JBI, of more conventional display technologies that will be commercially driven to smaller size, lighter weight, and lower cost.

4.2.3.2 Tailoring

The following technologies may greatly impact the goal of delivering the right information at the right time. The ability to tailor the presentation to a variety of aspects could significantly simplify user workflow by mining the useful information out of what is available. Tailoring relies on the ability to automatically capture user preferences as well as the workflow of the job in which the user is engaged. Data can be captured by developing and refining models of user tasks during training or in the early phases of a crisis.

4.2.3.2.1 Task and User Modeling

A number of techniques can capture both the user and task models. These models are developed offline and are then applied during run time. Workflow management is based on some of this technology. One issue is the robustness of these models in representing the task execution

process during operations. Several programs (Command Post of the Future, for example) have defined planned experiments to evaluate the potential performance gains that this technology might provide.

4.2.3.2.2 Information Needs Models

An augmentation to the development of task or process models is the capture of associated information requirements of the users of the JBI. This information can be captured as templates prior to a crisis, during exercise support, or during other preparation activities. The information is then tailored to a particular crisis and tuned during the crisis. This piece of the context could provide automatic subscriptions in the future.

4.2.3.2.3 Dialog Management

Another technique that can provide additional contextual cues for tailoring is dialog management. This area grew out of speech interaction work, which recognized that by bringing in the context associated with previous queries, commands could be generated to help reduce errors within the current interaction. Also, it enables the user to refer to a previous statement, which helps reduce the amount of input the user needs to provide. Dialog management focuses mostly on the problem of one person communicating to a single computer, but some early work was done on dealing with group dialogs, with multiple users, and with multiple agents.

4.2.4 Context Understanding

Context understanding is a technology that will enable the JBI to track its processing context. It supports the tracking of people in the command center (their location), is aware of their roles on the staff team, and monitors their current activities (for example, working, resting, or in a meeting).

Context maintenance is an important function for the JBI since many of the HCI interactions that occur will be either incomplete or ambiguous when interpreted in isolation. In some cases, components such as language understanding will use contextual information to disambiguate the user's statements.

As an example, consider the variety of contexts that a JFACC faces in the development of a single air mission for an air tasking order (ATO). It is the JFACC's responsibility to simultaneously discern the merits of the current air campaign strategy, the changing threat posture of enemy forces in response to that strategy, the readiness and availability of aircraft and crews, and the ability of the logistics tail to support alternative air mission scenarios. Each concern and assessment shares data, the significance of which changes with each concern or context. A logistics failure might be manageable for parts of a mission but critical for alternative strategies. The depth of concrete on tarmac might be fine for housing and emergency shelter but catastrophic for aircraft beddown. The environment of the JBI must offer the data for the user in a form that conveys the meaning of the data in the context of operation. Thus, the JBI, with knowledge that the JFACC is considering emergency housing options might offer the information that a small remote airfield is suitable.

4.2.5 Collaboration

The final class of interaction technologies to be addressed is collaboration. These technologies support human-to-human collaboration, human-to-agent collaboration, and human-agent-human collaboration throughout the JBI. The assumed purpose of this collaboration is to support problem solving and decision making. The most common commercial form of computer collaboration involves the use of video teleconferencing and shared “whiteboard.” However, collaboration software is maturing swiftly allowing users to share applications and other more complex software environments.

4.2.6 Command Summary

These interact technologies enable improved decision making (the goal of command) by (1) ensuring tailored, natural, and timely delivery of information; (2) the ability to rapidly and selectively drill down in areas of concern to improve understanding of the situation; (3) the ability to rapidly and easily collaborate with respected colleagues and intelligent agents to evaluate evidence about the situation and possible courses of action. Thus, the commander is better informed, has enhanced situational awareness, and thus is in a position to make better decisions.

4.3 Planning

The JBI users illustrated in this section are distributed inside permanent buildings with full power and communications connections. Their tasks are to develop a plan to destroy the TEL. The JBI supports collaborative problem solving.

This includes the requirement to determine availability and readiness of warfighting assets so that alternative Courses-of-Action (COAs) can be formulated and evaluated. The commander’s staff queries the JBI to determine the status of potential assets, and once determined, to establish rules for these assets to “subscribe and publish” data in accordance with the requirements levied by the commander’s “warning order” or course of action. Asset commanders are then directed to *subscribe*. A weather specialist *publishes* weather status (see Figure 7).



Figure 7. Collaborative Problem Solving

4.3.1 Data Capture

Concepts for capturing data are closely tied to the JBI operational concept of “publish and subscribe” for the collection and dissemination of data among its participants. To the extent possible, these functions should be automated but care must be taken not to overlook the vital function of human oversight for filtering unneeded and inaccurate data prior to dissemination. Also, the concept of data capture is extended to include data elements that have, until now been overlooked as needed within a warfighting operational environment. These include capture of speech (verbal command directives, meeting transcripts, etc.), gestures (pointing, nods, etc.) and possibly control actions (keyboard entries, weapon release sequences, wheels-up indications, etc.).

Currently, ASR technologies apply to data capture requirements of the JBI in two major categories:

1. Dictation services that enable creation of source data such as intelligence reports, battle damage assessment, commander’s intent, etc.
2. Computer telephone services that enable users to interact with computers over standard telephone channels.

4.3.1.1 Template-Driven Input

Templates are an effective method for providing “standard” defaults to the data field to aid the user through the application initialization process and workflow definitions. Smart “wizards” and “active templates” exploit Artificial Intelligence and software “agent” technology to aid the user through complex decision trees and can be used to monitor the data flow to alert the user to important changes in the data and/or requirements to focus on specific task sequences. It is anticipated that the JBI will heavily rely on agent-based Active Template technology (currently being developed by DARPA) to help users instantiate JBI environments, monitor and maintain data and workflow, and alert users to critical causal and temporal dependency conflicts.

For example, the Active Templates program will produce a robust, lightweight software technology for aiding in the automation of detailed planning and execution for military operations and other highly coordinated activities using a plan spreadsheet metaphor. These templates are distributed data structures whose variables will be linked to live data feeds or problem-solving methods. Thus, Active Templates will assist with automated planning and execution by capturing, improving and updating critical information such as current state, goals, constraints, alternative actions, standard defaults, problem-solving context over time, decisions, and rationale.

4.3.1.2 Smart Room Devices

The notion of a smart room is one wherein embedded sensors and processors observe the participants and infer various kinds of data input from what it sees. For example, if a person makes a gesture toward a screen across a room, rather than rely on special display or pointer technology, a camera (or other sensor) would recognize that the gesture applies to that display. This approach creates a totally unencumbered interaction with the user. One example of this type of technology is an MIT project called the Intelligent Room. This project exploits embedding

computers in ordinary environments so that people can interact with them the way they do with other people, by speech, gesture, movement, affect, and context. Using similar technology, a command staff can interact with each other and also the JBI seamlessly.

4.3.1.3 Annotation

This technology can be either text based or speech based. The idea is to capture analysis and conclusions that need to be attached to a location within a document. There are technology issues on how to store the annotations as well as how to display the annotation. This technology is important for both creating very convenient mechanisms for individual capture, but also for supporting asynchronous collaboration mechanisms.

Examples of annotation within the JBI environment are establishing a target list or creating an ATO for an area containing critical infrastructure. In creating the ATO, the JFACC may need to view map and reconnaissance image data to confirm the coordinates recommended for strike. Once confirmed he or she might select aircraft and weapon type as well as strike time, combat support, or other related components of this small piece of the larger operation. It is envisioned that the JBI would exploit natural language together with gesture processing to enable the JFACC to speak these data and provide context by gesture. He or she might command: "Strike this (gesture) bridge on day 22 at 0925 with F-16 armed with GBU, two sorties separated by 5 minutes." The data elements of this command are then "attached" to the coordinate location of the target site as designated by the gesture action. In this way the map becomes an active repository of target and air campaign planning data through the annotation process.

Annotation can also be used to reduce clutter in complex data-rich presentation. For example, the details that relate to the strike mission might be viewable only when a cursor is in a certain region or when the speaker is a certain person. This facilitates the process of "drill down."

4.3.2 Presentation

Presentation technologies include data visualization, holographic and 3-D displays, large screen displays, and active templates to tailor the information to user needs and preferences.

4.3.2.1 Data Visualization

It is clear that data visualization will play an important role in information comprehension within the JBI environment. Also, enabling the user to interact directly with the data, for example to drill down or query, greatly enhances his or her ability to deal with data complexities in dynamic environments. There are significant technologies in the development of basic visualization primitives to provide high value presentations that are individually crafted by specialist's technologies (for example the System for Automatic Graphic Expression [SAGE] from CMU). Also, there is emerging automatic display generation that provides visualization constructs that meet the requirements of users' cognitive processes and workflow. Creating environments that provide both high value information visualizations combined with natural interaction should improve a user's ability to quickly find the right information with the right supporting data.

A good example of this presentation technology is shown in Figure 8, which depicts details of Napoleon's campaign against Russia. This presentation was automatically generated from a file

that contained data relating to his force size and location, the air temperature that impeded his assault, and the geographic area of the campaign over time. The complex interrelationships between the various data elements are clearly presented enabling easier interpretation of the campaign.




Figure 8. History of Napoleon's campaign against Russia (five dimensions displayed in a single graphic: location, temperature, size, battle location, name, and date)

4.3.2.2 Holographic/3-D displays

Since much of what the user is going to interact with is geospatial, it makes sense that in particular situations, the most effective way to present the information is in a true 3-D environment. However, 3-D visualization techniques are also proving to be of significant value in areas of clustering and data fusion where the high number of dimensions associated with the data dictates the need to provide user views from different perspectives and projections. Currently, 3-D visualization techniques have the disadvantage of requiring specialized glasses or environments, but research is gradually reducing the encumbered nature of this interaction.

4.3.2.3 Large Screen Displays

While it is anticipated that the future battle commander will be highly mobile, there is also a need to be able to view all relevant information with clarity, speed, interactivity, and organization. To date, the data detail for display systems has been one of the bottlenecks in the information channel to the user. With Large Screen Display technology such as Sarnoff's System Technology for Advanced Resolution (STAR) or the AFRL Data Wall, there will exist a scalable, interactive display technology that will support very large display surfaces (20ft) with hundreds of mega-pixels. This will enable collaborative users to interact with maps and other information displays in a very natural manner. The displays could be used for small group problem solving or small or large group briefings. A key part of the notion is that a workstation is not tied to a particular element of the screen and that displays could be arbitrarily positioned.

4.3.2.4 Active Templates

Active Templates (being developed by DARPA) is best thought of as an “intelligent spreadsheet” for planning, information monitoring, and execution re-planning. A user-defined artificial intelligence (AI)-based automation layer that enables definition of logical relationships between the data elements of the spreadsheet provides the capability for software processes to prefill and check for temporal and causal consistencies between the elements. Automation can be implemented in a variety of ways, including

1. Simple scripts attached to plan slot values
2. Data sentinels that trigger on some condition on the plan or execution data
3. Rules from which new information is inferred
4. Constraints that propagate within the current template and across the network to others
5. Links that signify an implicit relationship that the system cannot reason about, but through link analysis, can identify at least the relationship and remind the human user to act
6. External calls to a generative planner, scheduler, or machine learning algorithm once sufficient information is available
7. Calls to simple or sophisticated autonomous agents that can broker or handle difficult problem solving tasks on behalf of the user

Automation is meant to solve problems by tailoring default templates, merging them with other templates to handle dependencies, and reacting in appropriate ways to new information, even when it is incomplete or possibly wrong.

Active template technology enables some of the important notional concepts of “publish” and “subscribe,” which are underlying operational themes for the JBI. If active templates are created to allow collaboration (see Section 4.3.3) among users, then the templates can be used to manipulate distributed data fields that are published to various subscribers. The publishing and maintenance of these data can be performed through both human and intelligent software agent processes.

Table 1. Example of an Active Template. When the location is entered (in the right-hand column), the generic values change to tailored values for that location.
(Question marks identify information requested from the user.)

Airfield Seizure Template	Generic	Tailored
Location		Redondo Airfield
Int Staging Base	???/300km	Aviano/353 km
Ingress	???/???	C-141/2
Debarcation Method	???	Airdrop
Task Force	Bravo	Bravo
Noncombatants/ROE	???	Civ C2/Mosque
Action Sequence:		
R&S	Airdrop/???/???/8	Airdrop/10km/foot/8
Initial Strike	Comm hit/Tower sz	Comm hit/Tower sz
Tact Threat Neutralization	Patrol hit	Garrison/HQ/Patrol
Tact Objective	Clear runway	+establish air C ²
Initial security	Establish perimeter	+Secure fortification
Continuing Defense	Hold 6 hrs	Hold 8 hrs
Egress	C-141/???	C-141/2

In Table 1, an example template is shown that includes the data fields associated with an operational plan for securing an airfield. A generic template provides the necessary elements of the checklist for carrying out the operation. Associated with each cell in the checklist is a logic program or mathematical procedure that helps the user determine appropriate values for the data fields. Thus, for example, when the user types in the fact that the Redondo Airfield is the target location, the template responds automatically with default data fields for the intermediate staging base, method of debarkation, and type of transportation and logistics support.

4.3.3 Collaboration

The JBI envisions a collaborative ATO generation capability using Active Template spreadsheets that are shared between the participants. Thus, rather than just sharing the same view or providing annotation to data, the ATO Active Template would provide the application to each user. When data fields are added or modified by one planner to designate the change in tail number or mission objective, each participant in the collaboration would share this action and observe the change in the data fields automatically. Sharable collaborative applications such as these offer tremendous potential for reducing the development time for detailed operational plans. Individual users can work on several parts of the plan simultaneously, while the system keeps the implications and dependencies of operational decisions visible to the planning group as a whole.

4.3.3.1 Advanced Whiteboarding

Collaborative “whiteboards” enable several users to share and visualize data simultaneously. A method for editing these data in real time, either sequentially or simultaneously, is usually supported. The most common military application involves the shared use of map data and operational plans. Technology for providing advance whiteboard services is moving swiftly in

the commercial marketplace, and large multi-user collaborated workspaces that support simultaneous data manipulation and update are now supported. It is clear that this technology has significant payoff for the JBI environment, particularly in the area of plan/re-plan generation and situation assessment.

4.3.3.2 Mixed Initiative

Humans and computers have complementary skills. While humans are good at conceptualizing meaning and understanding the context of goals and objectives, computers are good at manipulating data and managing details. While humans are good at applying intuition and strategic knowledge, computers are very good at executing quantitative analysis methods. The goal of mixed initiative planning is to create a seamless application environment that exploits the relative strength of each domain and thus provides better services than could be expected using each independently. Investments in mixed-initiative planning technology DARPA have yielded impressive results in the area of resource management and scheduling. There has also been initial application of this technology to the broader problem of strategic planning. For example, a planner within the JBI must constantly assess whether the current operational course of action will achieve the operational objective. In the Planning and Decision Aids (P&DA) Program, DARPA was able to demonstrate the use of computers for quickly manipulating data to calculate tactical moves in response to changes in strategy dictated by human planners. As is the case in the game of chess, computers are well suited for fast calculation of optimal moves based on tactical doctrine that remains largely invariant. The P&DA program demonstrated the very rapid generation of air logistics support plans in response to changing air combat environments.

4.3.3.3 Group Interaction Devices

One of the more interesting opportunities for visualization technologies involves solving the problem of making it intuitive for large groups to share complex data within a collaboration space. Two emerging technologies are the use of 3-D displays and data walls that support the capability of allowing multiple people to interact simultaneously.

4.3.3.3.1 3-D Sand Table

One interaction device whose value is yet to be proven experimentally in the command center is the 3-D sand table. This technology allows a small group of people to interact directly with a 3-D view of terrain and units. An example of this is the Dragon system at the Naval Research Laboratory, which was one of the first examples of a Virtual Reality Responsive Workbench. There are a number of large screen display systems now that support group interaction with 3-D views. They all require special glasses to interact with them, and display screens can either be vertical, horizontal, or tilted.

4.3.3.3.2 Data Wall

One interesting ergonomic aspect of large screen displays involves the use of small groups interacting with a large “data wall” made possible by the large display real estate. With today’s display systems, there is only one mouse or pointer that a group would share. In the future, it is

envisioned that several groups using several pointers and editing devices will be able to simultaneously interact with very large data sets on the same display.

4.3.4 Planning Summary

Group interact technologies support effective and efficient planning by identifying the information that should be considered (template-driven input), supporting sophisticated interactions with the data through annotation and data visualization methods, and enabling collaboration across distances and time.

4.4 Execution

There are three different types of JBI users illustrated in this section: (1) special operations personnel visually verifying the TEL location and status and transmitting this information over secure telephone lines, (2) Army support personnel overseeing the clearing of the attack zone and entering the status using personal computing devices, and (3) an F-18E pilot assigned to strike the TEL. The illustrations show how the JBI supports automatic formatting and filtering of information for users in different physical environments.

The selection of a course of action is a collaborative one involving the commander, his or her Information staff, senior leadership of the joint staff, the task force commander (if assigned) and selected component commanders. Also, detailed planning and analysis will be expedited and made more accurate through collaborative applications that allow plan developers to jointly create complex and interoperable plan segments that play together. As shown in Figure 9, the staff *subscribes* to collaborator assessments and *publishes* ATO execution. The JBI performs automatic formatting and filtering. The JBI environment will allow the commander to dynamically generate and modify business rules that reflect the commanders' tasks, desired preferences, and security restrictions that must be enforced. At all times the commander will have total control of the dissemination of data and the establishment of work-group environments that reflect multilevel security, and need-to-know access controls that must exist in all military operation actions. In this chapter's example, the commander *controls* dissemination by granting SOF access rights. Data are *transformed* to meet device, task, and user requirements whether in the field or in the cockpit (see Figure 9).



Figure 9. Automatic formatting and filtering of information for field and cockpit

The JBI environment will provide services that translate published data into the proper format for the subscriber's requirements and interface modality. Thus, combat troops in the field might be delivered information about the timing and location of Close Air Support (CAS) events via heads-up display while the JFACC might view the sequence of air missions via a laptop. Also, in the example, when Special Operations Forces (SOF) publish the current location of enemy transporter-erector-launchers, these coordinates are received and translated to the appropriate data format to be seamlessly delivered to the cockpit weapon's platform aircraft that have "subscribed" to these events.

4.4.1 Data Capture

Since JBI users will find themselves in a wide range of environments, the system will need to be able to support capture from a number of devices referred to as personal computing devices. These processing devices span the range from palm top computers to wearable interfaces.

Small computers, ranging from laptops to palm tops, to pen-based computers are impressive in terms of computing power and display capability. With laptops, one pays a price in terms of power and weight and most laptops still require much keyboard entry. Palm tops, like the Toshiba Libretto, are certainly a lot less bulky, but they are very much disadvantaged in terms of display functionality and resolution. Pen-based computers, like the Fujitsu 2300, occupy an interesting middle ground, in that they provide better display capability than the palm tops, have the potential for much more natural interaction with the pen, and are less bulky than a typical laptop. The biggest problem is that the operating systems are not quite ready for pen-based computing. Ultimately, the JBI will have users with all of these classes of computers, and there must be facilities to support the wide range of input and output devices.

Within the JBI, laptops and workstations will, for the most part, be prevalent at command centers and detachments where intensive planning or analysis activities require constant evaluation of all data at one location. However, these computers are too cumbersome for the operations of mobile foot soldiers or SOF forces. Such forces require agility afforded by smaller, lighter weight packages. Furthermore, as forces strive for increased mobility and decentralization of command function, there will be an increased need within the JBI for small personal assistants at all levels of the command echelon that can provide seamless access to data automatically tailored to the user's needs and personal digital assistant format.

One type of personal computing device is a wearable computer, a technology that integrates computing devices with clothing. Typically, computing devices such as the CPU and batteries are incorporated into a belt; the input devices are potentially speech and some sort of tactile input; speech and helmet-mounted displays or glasses are the output. For many tasks where you need both hands to operate effectively such as maintenance, this may be the only way to effectively interact with the JBI. The DARPA Communicator program is another example of a project that is pursuing effective computer interactions without the use of traditional screens or keyboards. The Communicator program is not only pushing harder on speech technology but recognizes that effective computer interaction will require new techniques in dialog management.

4.4.2 Presentation

In this case presentation includes Virtual Retinal Displays (VRD), nontraditional sensory stimulation, 3-D audio, and domain-specific workflow management.

4.4.2.1 Virtual Retinal Displays

The extreme presentation challenge is confronted when absolute mobility and hands-free interaction is required in high noise environments. While this confluence of conditions will rarely be encountered, they mandate the need for personal display devices. A number of maturing technologies exist that offer good display resolution with minimal hardware footprint requirement. Wearable computer display devices are viewed as viable technology options affording needed user mobility, but these devices must be small and easily usable. A highly advanced display technology that promises to provide this is the VRD from the University of Washington Human Interface Technology lab as well as other research organizations. VRD is based on the concept of scanning an image directly on the retina of the viewer's eye. This technology can result in either head-worn or handheld displays with reasonable resolution, can allow the user to view the "real" scene simultaneously with computer-generated data, and provide needed mobility.

The use of retinal imaging devices is also being evaluated for use in helmet mounted heads-up displays for cockpit targeting applications as well as for presenting enemy force position to ground force soldiers in close combat situations. In the future, VRD technology may replace head-mounted goggles and special glasses for 3-D and virtual reality environments.

Within the JBI, a "networked soldier" on the ground could communicate directly with command and subordinate forces through the Remote Piloted Vehicle (RPV)-based digital cellular network.

Data transfer from the soldier would be through natural speech so as to enable maximum mobility while communicating. Global Positioning System (GPS) and other measurement sensors he or she is wearing provide continuous data “publication” for monitoring position, health, and status. Command could pass or provide “subscription” service, images, and data to the foot soldier through the network to a VRD eyepiece providing instant status of enemy or coalition force components without encumbering his or her ability to move and fight. VRD eyepieces are unique from related miniature television screen eyepieces, in that the user can simultaneously view the image and view “through” the image and thus are able to maintain a natural perspective of his or her surroundings.

4.4.2.2 Nontraditional Senses

The realism to be offered within virtual reality environments is likely to dramatically increase in the next decade, offering greatly enhanced human-to-human interface for geographically separated collaborators. Olfactory cueing could be used to signal presence of toxic waste, fire, or smoke. Tactile prompting can indicate direction, as applied in the tactile vest, or alerting, as a poking in the shoulder. Taste can also be used to alert soldiers in the field; specifically, dental mounted sensors can release different flavors to indicate enemy actions.

4.4.2.3 3-D Audio

3-D audio provides an interaction technique that has application in either a group environment or for an individual when the directionality of the data can be exploited to focus attention or increase information content. 3-D audio has been used to: (1) improve intelligibility, (2) provide navigation cues, (3) warn of threats, (4) support targeting, (5) indicate location of wingman, (6) give location cues to Air Traffic Controllers, (7) help the blind navigate, and (8) hands-free communication. Also, 3-D audio is of significant value within virtual reality environments where group interactions are more naturally induced.

4.4.2.4 Domain-Specific Workflow Management

Decision support tools that provide support for group workflow processes are also emerging in the commercial marketplace. These technologies offer the capability of decomposing complex multitask processes and distributing them over a dynamic set of execution assets. In general, workflow management techniques for large groups can be difficult especially when priorities and task failure penalties vary greatly, as in a military operational environment. Current technology focuses on the development of operator models that can be used to assign assets based on the prediction of future states. These techniques offer the hope of providing automated support for group activities, such as planning an operation, that have a reasonably well-defined process. The challenge usually surfaces when the modeled process does not closely match the process required for execution.

4.4.3 Collaboration

Collaboration involves many forms of communication among a group of users. Sharing of data, either synchronously or asynchronously, can greatly enhance this communication or thereby positively influence rapid decision making. The JBI will exploit collaboration technology at both the data and application layers. Data collaboration will involve sharing files, images, maps,

plans, and other workspace objects that support continuous operations. For example, sharing logistic ammunition inventory data with multiple operational planners involved in generating ATOs will greatly improve the efficiency and speed of the ATO generation process. Sharing will also be supported at the application layer. For example the generation of the ATO itself may involve the use of a shareable ATO application where several operational planners are simultaneously generating different parts of the common plan.

4.4.4 Execution Summary

Execution of the mission is supported by interact technologies that automatically capture data, present information in forms tailored to the user and device, and support collaboration among various device formats.

4.5 Combat Support

There are seven JBI users illustrated in this section: (1) a crew chief inspecting the F-18E aircraft selected to strike the TEL, identifying a failed line replaceable unit (LRU), and entering the status using a wearable computer, (2) logistics personnel at distributed warehouses fulfilling supply requests using personal computers, (3) load masters overseeing loading and unloading of supplies onto transport aircraft using palm tops, (4) F-18E schedulers aboard ship with limited communications bandwidth, (5) air refueling schedulers at the forward air operating centers in tents with limited power and communications, (6) carrier deck personnel launching and retrieving aircraft and recording launch times with pal top computers, and (7) pilots in the KC-10 and the F-18E aircraft rendezvousing and refueling. The JBI supports automatic data capture.

While most of the example described in Sections 4.1 through 4.6 focuses on information structured around the conduct of combat operations, it is equally important to emphasize that effective operations can be assured only if there is effective Combat Support (CS) and Combat Service Support (CSS). The JBI environment contributes greatly to the concept of a tightly coupled operational and logistics system. Logistics failures can be published in a timely manner to create operational alerts in the form of shortfalls that can be better filled through subscriber applications and system services that can search for available inventory or re-route supply lines if needed. In one example, a crew chief identifies a failed LRU on the aircraft selected for the strike against the TEL. This is automatically captured and *published* to JBI. Logistics *queries* JBI to locate a functioning LRU. The loadmaster captures LRU arrival using the bar code. A crew chief installs the functioning LRU. The F/A-18E scheduler *subscribes* to aircraft status and assigns the sortie. The F-18E launches (see Figure 10).



Figure 10. *Automatic data capture for combat support*

Through the JBI environment, users at every level of the operational echelon have abilities to publish all operationally relevant information. If the information is important to many users, their subscriptions ensure that they all receive the data in a timely manner. At the same time, the JBI ensures that data translators and filters protect the subscriber from “data overload” and incompatibilities with operational setting or hardware presentation facilities, thus meeting the JBI goal of providing the right information, to the right person, at the right time, in the right media and right language, and at the right level of detail.

4.5.1 Data Capture

Data capture includes both information extraction and point-of-use capture through a variety of devices.

4.5.1.1 Information Extraction

Information extraction (IE) is the ability to identify and capture useful information in text and store it in a structured form such as database records. IE capabilities are typically divided into two levels based on the complexity of the information they extract. Shallow extraction refers to extraction of simpler types of information such as entities (the names of people, facilities, locations, etc.), numerical information (monetary values, percentages, etc.), and simple events (action = verb). Deep extraction refers to extraction of much more difficult types of information such as complex events (also known as scenario characterization). The ability to capture and transform data (free text) into structured information is useful in many different ways:

1. Automate the processing of very large volumes of free-form text, saving time and labor
2. Enable persistent storage of the extracted, labeled information in databases
3. Enable information to be used by other user support tools (for example, analysis and visualization tools)

There are two basic approaches to developing IE systems: (1) the knowledge engineering (rule-based) approach and (2) the learning (statistical) approach.

IE has obvious implications for data input and information capture for a JBI. It may also provide novel methods for query formation, and when coupled with a graphical user interface, IE can facilitate understanding through information visualization.

4.5.1.2 Point-of-Use Data Capture

Point-of-use data capture is critical to the information support and administrative portion of the JBI. Data are generated in a great variety of locations, but they are not always captured at the source, or in a usable form to be shared with those who need them. Examples include administrative data such as health care and financial records, maintenance data that are collected during maintenance inspections or that are generated by embedded information systems in an aircraft, and data that show the consumption of supplies such as fuel and ammunition. With the emerging availability of wearable computers and wireless technology, it is reasonable to expect to capture much of these data at the point where they are generated and automatically make them available to the JBI. The challenge is not so much with the technology itself, but with developing the business rules and processes that define the automated data capture process.

4.5.1.2.1 Bar Code

Bar code technology is one means to minimize human data input errors. This technology can be a very effective and reliable means for users to identify the item they are working on and to automatically report status and parametric input data. Thus, rather than burden the user with the requirement to enter significant quantities of data about an item of interest, one swipe of a bar code coupled with a data base provides the data automatically.

4.2.1.2.2 Smart Tags

Advancement in bar code technology is a class of devices referred to as Automated Identification Tags (AITs) or “smart tags.” These have the advantage of being able to be queried, either actively or passively, from a distance, and many technologies support significant volumes of dynamically loaded data. So these tags can contain data that would have to be accessed from a database. While the cost for the most capable active tags is still relatively high, technology is moving rapidly to reduce cost, expand on-tag data storage, and greatly increase interrogation ranges for both active and passive methods. Also, AIT devices are becoming functionally more complex. The GPS locator, for example, has the ability to not only tell about itself, but also to provide its own location.

There are many applications of AIT coupled with GPS technology that directly benefit the JBI. Two of the most significant ones are in the areas of asset visibility and readiness condition. Cargo and equipment are able to “report,” through active or passive interrogators, their identity

and location while dynamically being repositioned in the theater. At the same time, diagnostic or autonomic measurement sensors could update the status of the cargo while in transit, providing data to human or automated planning algorithms whose job it is to maintain high combat service support. Smart tags can also be thought of as hardware agents sending messages, providing inventory, requesting parts, or alerting an end user that the cargo has been spoiled. These data are published to the JBI and updated automatically.

4.5.2 Presentation

Haptic (force-reflecting) interfaces can provide useful kinesthetic information in virtual and real environments. Haptic devices, for example, can use magnetic levitation to provide the perception of physical interaction with simulated objects and environments on computer screens. These devices are unique because they enable people to not only touch objects, but to reach in and manipulate them in three dimensions as well. Several haptic interfaces are currently in use including data gloves, virtual switch controls, and four-degree-of-freedom controllers used to train doctors in surgical procedures. Another specific example of a haptic interface is the tactile vest that can be used to provide an alternative method for cueing the user to look in a particular direction by simulating a tap on the shoulder or another kinesthetic cue. Haptic devices may enjoy an important role in virtual environments but can also be of value for augmented sensory feedback in traditional environments as well.

4.5.3 Collaboration

In any meeting, there is usually a facilitator who makes sure the right things get done and the right interactions occur. There are some advanced technologies being explored that would act as facilitators in shared collaborative computer environments with the use of software agents. Thus, a small group that is using collaboration technology to accomplish a particular task could use a facilitation agent to expedite the process.

4.5.4 Combat Support Summary

Interact technologies automatically capture data at the point of use ensuring the greatest accuracy and currency, extracting information from the captured data, and presenting it in task-specific formats across all echelons of combat support.

4.6 Key Interaction Technologies

The interaction technologies discussed in this chapter are designed to improve the productivity of the JBI users by providing automatic data capture, user-centered presentations, and collaborations among users. Given the information-centric nature of the JBI, it is anticipated that much of the technology will focus on the need for users to effectively and efficiently interact with computer display/presentation systems and databases. Also envisioned is the need for the system to anticipate users' information needs; instigate the search for relevant data; push, update, and maintain the currency of these data; create information from the data; present the information in the form tailored to the user and his or her current needs; enable collaboration across distances and time; and maintain pedigree for drill down if required.

Thus, it is expected that the JBI will demand automation in many phases of the interaction. Furthermore, each user of the JBI may have different time-sensitive priorities for information, and automated interact functions should be tailored to meet specific user needs. The automation of specific functions will continuously evolve; taking advantage of lessons learned from exercises and real-world crises. The following identify the three primary roles—automatic data capture, user-centered presentations, and collaborations—for interaction within the JBI.

4.6.1 Automatic Data Capture (Including Context Filtering and Formatting)

Once instantiated, the JBI will exploit and maintain a significant volume of data, the value of which will vary depending on time, context, and user workflow requirements. To be able to deliver the right information to a user, the JBI needs to provide tools to ensure that it has captured the appropriate input data throughout the system, employed context filters that are tailored to the work-flow processes of users, and translated these data into formats consistent with user display and decision-making processes. Traditionally, attempts to capture relevant input data have put a significant administrative burden on both operational and support staffs, often introducing unacceptable latency and resulting in incomplete and inaccurate data. Within the JBI, there is a need for a wide range of automated techniques for capturing information, routing it based on subscriber needs, and maintaining its pedigree for authentication and ambiguity resolution.

Although still fairly primitive, several hardware and software techniques exist for automatically capturing data for information systems. Hardware examples include the use of bar code readers (primarily for logistics applications), the use of wearable computers with voice input and wireless communications to capture data in the field as a byproduct of other activities (such as maintenance), and the automated extraction of data directly from deployed sensors. Also, technologies are beginning to mature in the use of intelligent software to search very large data sources, retrieve context-relevant text documents or segments, and return them to the user for analysis. Included in this latter category is the automatic extraction of data from messages, images, intelligence reports, and even briefings and meetings. Extraneous information that falls out of context with the user's task must be filtered so as not to overburden the user with unnecessary data. Automation will be used to develop task and user profiles, which will guide data search and extraction and presentation of information.

Once found, the information provided to a user must be in the right format and tailored to the specific task the user is performing. It should also be presented in a format and style that is best tailored to the specific user workflow. Automation will be used to develop task and user profiles, which will guide the presentation of information. The devices available to a user will dictate how information should be formatted. For example, narrow band communications on the battlefield or devices without a display (such as a telephone) will constrain how the information (such as the content and information) can be presented. Related to this constraint is the critical need to ensure that time-critical or priority data are delivered in a timely and predictable manner.

Data will also be captured from the user as it is created in the natural process of dialog between humans and/or humans and machines. Natural language translation programs are maturing

rapidly and can now capture conversational speech and identify or label the speaker. While the performance of these applications is still poor, commercial interest in this technology is expected to provide impetus for intense research. Likewise, video-capture and gesture identification systems coupled with natural language interpretation systems enable the inference of context from the speaker's words. Pointing at a specific location within a room or on a display enables better understanding and removes ambiguities from the spoken language. Gesturing while speaking also enables the seamless data entry of facts that originate within a dialog between people. For example, the speaker is able to interact with the data repositories using commands such as "Get me the latest intelligence estimates for this mountain region here," or "The underground chemical plant is probably right there." "Here" and "there" can be automatically translated into the georeferenced coordinates of the speaker's gesture point.

The process of annotating figures and text is also greatly facilitated through combined gesture and speech interpretation. The CMU Quick Target system enables image analysts to view reconnaissance photographs and quickly provide annotation and text analysis to the specific area of interest. The analyst's words are captured and translated into text and become an object in the target or intelligence database associated with the specific coordinates of the map image. This kind of human and machine interaction is most efficient and provides a foundation for creating and sharing knowledge between several users working on the same problem. While voice and gesture interpretation technology is most prevalent in the setting of the "command center" or analyst's workstation, this technology can also be used with small personal assistant devices where the mode of gesture is limited to finger pressure or pen tip.

4.6.2 Task-Centric Presentation

The understanding of a situation and the available options depend critically on presenting the information in an appropriate form. Presentation format will exploit multimodal sensor input from a combination of visual, aural, and haptic interfaces. The use of 3-D graphics, geospatially referenced presentations, animation, image zoom, and moving forward and backward in time are some of the tools that are available. But the presentations must be tailored to the workflow task and to the preferences of a particular user. What is presented in the cockpit may be very different from what is presented in a command center.

For example, the University of Southern Florida has developed a method for presenting aircraft flight status information using geometric shapes that take advantage of the operator's peripheral vision. Rather than viewing altimeters, horizon, air-speed indicators, etc., the pilot is presented with a geometric view of these measurements all within one presentation format that enable him or her to fly an aircraft within minutes of being trained. This same research center is also developing a haptic vest device that allows a user to feel pressure in different regions of the midriff and back areas of the body. The pilot can use the locations of these pressure points to infer direction. Thus, the pilot can be given data on target bearing and rate of bearing without looking at a conventional cockpit instrument.

In addition to format, the amount of data and their "flow" (update rate) as presented to the user should be the minimal needed to maintain an accurate status of the situation. Very large display

screen devices that hold large amounts of detailed data can be used to provide context between different “types” of operational information when viewed from a distance. For example, a large area, high-resolution display panel may contain details of the supply status (3-D bar graph) for critical items on a weapons system for an Inventory Control Point (ICP). At the same time, it may contain the location for all weapon systems and their readiness status with reference to waiting parts required to gain or maintain full readiness. Large screen, high-resolution displays enable these data to be viewed in aggregated form from a distance, enabling cooperative planners to quickly assess the status of supply as it may affect the status of readiness.

Large screen displays can be thought of as collaborative devices in that they bring users together to view the same data while focusing on different problems. In addition to offering large amounts of display real estate, they enable users to be physically together during the decision making process. They are popular in the command centers where the convention command structure dictates the collaboration protocol.

4.6.3 Collaborative Problem Solving

The quality and timeliness of decisions are dependent on a shared understanding of the problem and the available options. A representative problem is the production of ATO, which involves intelligence, operations, and logistics participation. In the past this has been a serial process and, therefore, a slow one. With split-base operations it also involves separation in space. With 24 by 7 operations there is also asynchronous collaboration where information is shared across the shifts. The collaborative problem-solving capabilities in the JBI will enable teams to collaborate within a shared information environment to speed the problem-solving by working in parallel rather than serially and to work with partial solutions when there is still flexibility to adapt. In addition to the collaboration between individuals and groups, there will be collaboration between people and software agents.

JBIs collaborations can be enhanced through the use of high-resolution hardware and advanced graphics techniques. For example, it has been shown by Sarnoff Laboratory (and others) that the use of 3-D immersion graphics and virtual reality displays greatly improve the decision makers’ ability to sort out large amounts of data associated with Air Combat Coordination operations. Rather than viewing icons for target type, velocity vector and altitude, the commander is able to view a 3-D picture of the operational scene and place him or herself anywhere so as to increase awareness of trouble spots and reduce ambiguities from the data. These techniques are sometimes called “3-D Sand Tables,” as they provide a method for several decision makers to participate together to better understand different aspects of a complex operational environment.

4.7 Interact Summary

This chapter provides a view into the range of interactions that will need to be supported across a wide range of users and also a wide range of environments, from groups of users in the command centers to individuals in the field or on the flight line. The illustrations in this chapter show how the JBI tailors information for the individual users depending on the physical environment in which they work and the computing equipment in use. A large number of relevant interaction

technologies are catalogued in Appendix D. The value of the JBI is realized when the users can quickly comprehend the information required for their tasks, hence this section details many promising user-computer interaction technologies. The value of having the right information available must be maintained through the use of a proper interface.

(This Page Intentionally Blank)

Chapter 5: The JBI Input Segment

5.0 Introduction

The previous chapter described the variety of ways the JBI supports user interaction. While much of the *interact* description centered on different display and output mechanisms, just as important are the multitude of input devices and mechanisms the JBI makes available. Moreover, behind these devices must be processes that make sense of users' inputs. The JBI must do this for user inputs as well as for inputs from other information systems.

This chapter describes the function and design of the JBI Input Segment. The first part of the chapter describes the variety of expected JBI inputs, some of which are shown in Figure 11. Next follows a description of the JBI's internal SCR, a schema describing all data stored in the JBI. The functionality of the JBI Input Segment, including a detailed discussion of fusion using agent technology, is then described.




Figure 11. The JBI Input Segment deals with inputs such as combat support products, fusion products, planning and execution products, command guidance, and user information products

5.1 Types of JBI Inputs

This section describes the envisioned inputs to the JBI. Some of these are shown in Figure 12. The JBI does not replace legacy applications; they will interface with the JBI through the use of wrappers, mediators, and agents. External sensors provide information to the JBI. The JBI may submit data to fusion engines to generate higher-level knowledge. The JBI tasks sensors to provide specific data. Using web-based interfaces, users direct the JBI to task sensors.




Figure 12. Inputs to the JBI

5.1.1 ISR Resources

The JBI will draw from throughout the ISR community. The potential inputs include databases; collection sensors and platforms; processing, exploitation, and dissemination systems; tasking and management systems; analysis centers; and units and agencies that support ISR activities. The JBI's relationship to ISR resources includes the following:

- The JBI must have current knowledge of the capabilities and availability of ISR resources down to the individual resource level (for example, not all U-2 systems have the same sensor version and/or packaging capability). Additionally, the JBI must have knowledge of resource ownership, procedures for acquiring resource support, and how to connect to resources for support.
- The JBI will query ISR resources to determine their suitability for supporting a given JBI application and will arrange, through appropriate channels, for specific support to that application.

The JBI will submit tasking, processing, exploitation, and dissemination (TPED) requests to the collection management system. The JBI may also query as to resource status and capability and TPED request status. The JBI will need frequent feedback from the collection management system as to the status of TPED requests, particularly changes in planned support.

5.1.2 Joint Planning Network (GCCS) Inputs

The JBI will make available the JTN product portfolio in its entirety. The JBI will use the integrated Joint Planning Network database service to obtain intelligence information that relates to ongoing strategies, plans, and assessments.

5.1.3 Joint Tactical Network (Link 16) Inputs

The JBI will bring in real-time battlefield track data.

5.1.4 Integrated Broadcast Service Inputs

The JBI will bring in real-time data that have been prefused.

5.1.5 Global Data Warehouse Interface

There are evolving concepts and programs for information systems (for example, Imagery Product Archive and Intelink) that will provide significant capabilities for the JBI to acquire information needed to support military planning, execution, and assessment. The JBI must be able to submit queries and information requests into these systems and receive status and data in return.

5.1.6 JFACC Inputs

There may be a large number of knowledge workers whose products will provide inputs to the JBI in accordance with the JBI business rules. These may include, for example, members of the JFACC operations and intelligence staff who are responsible for providing information support to primary JBI users. Other inputs may come from operations and intelligence personnel at higher, lower, and adjacent joint and component echelons, asset and/or resource managers at all echelons, and personnel in the intelligence community.

The JBI will contain JFACC system products relating to the military planning and execution processes. It is assumed that the “push” of these products through published business rules will be provided to the JBI interface system as a common capability.

The JBI may make queries about JFACC products (for example, status and numbers of planned activities) and expect a reply. JFACC systems will alert the JBI when a document element is added. Also, the document manager will advise the JBI when products with fixed information linkages are modified, and when pre-identified changes have occurred or thresholds have been met.

The JFACC intelligence staff are responsible for a number of intelligence functions (for example, discipline and all-source analysis and reporting, threat analysis and reporting, analytical assistance to electronic combat, target materials development, and collection management) that provide support to military operations. JBI interface system will function as an interface for these elements (tools and personnel) to link their products into the JBI plan monitoring and battlefield awareness processes.

5.1.7 Inputs From Other Command Echelons

It is assumed that JBI-like capabilities and processes will be conducted at many echelons within the C² system-of-systems architecture. The JBI must establish and maintain collaborative linkages with the elements responsible for these activities. This relationship includes the ability to request as well as to provide objective-based information support and to coordinate joint information-gathering actions.

5.1.8 JBI User Inputs

The relationship between the JBI and the user is two way. The user may issue commands to the JBI. These commands may require responses, initiate actions, or set event triggers through the

use of fuselets. Additionally, the user may enter data into the JBI via interaction modalities. The JBI may alert the user about events or conditions that have occurred and require user attention, or about activities taking place that require user participation.

5.2 Managing JBI Inputs

As shown in the previous section, the JBI may accept inputs from a wide variety of sources. Some sources will be weapon systems in the heat of battle. On the other extreme, some will be systems in the continental United States providing support via reachback. In all cases, the JBI must perform certain operations. Specifically, the JBI must

- Ensure the validity of data or information made available for JBI users.
- Perform smart filtering on data or information before it is input into the JBI. This function helps reduce redundant and irrelevant information.
- Transform input information into a SCR so that the JBI can work with it.

5.3 The JBI's Structured Common Representation

The SCR captures the current understanding of the military situation within the JBI. The JBI software has offered the content; the user has modified or approved of the content. The SCR is the product of many computer-person or computer-computer interactions. What is in this data structure?

- Hierarchies of various types, in which objects are linked to each other to represent their relationships in the hierarchy. For military forces, the most used link means “is a part of.” (An F-16 is a part of an air wing that is a part of available air wings that is a part of ... and so on). Hierarchies based on many other relationships are also possible.
- Collections or sets. Some objects of interest to the current understanding of the situation are not usefully grouped into hierarchies. These are kept in unstructured sets.

A force hierarchy of the SCR has many levels. The “upper” levels refer to battlespace things that are of interest at the higher levels of command. At the lower levels are things that are engaged in the battlespace (for example, a battalion or, lower still, a tank group or a tank). At every level, the evidence from sensors or data that support belief in an object (physical or conceptual) at that level is linked. Though the beliefs may be uncertain, this uncertainty is recorded quantitatively. The SCR contains no completely *ad hoc* beliefs. (Computer scientists and decision support system engineers will recognize portions of the SCR as a “belief network.”)

5.3.1 Ontology

The elements and structures of the SCR get their military “semantics” or meaning from the JBI's *ontology*. The ontology is the JBI's vocabulary. If there is not an ontology entry, the object or concept cannot enter into the SCR. Note that the ontology is part of the *knowledge* used by the fusion process and is not part of the process itself. The JBI's ontology—both development and use—is discussed in several sections of this report.

Specifically, it is envisioned that SCR will represent

- Hierarchies of forces
 - Blue
 - Land SCR
 - Marine SCR
 - Air SCR
 - Joint SCR (to the extent that it differs from the first four)
 - Coalition partners' SCR
 - Red
- Collections of other objects that are part of the theater war or sustainment situation but are not captured in a hierarchy
- Links
 - Between elements of the various hierarchies
 - To evidence (sensor or other data supporting belief in elements)
 - Between objects in collections and elements of the hierarchies, representing additional evidence, additional supporting data

5.3.2 Processes

Processes of the JBI supported by the SCR include

- Input to the visualization processes: the SCR is the primary information structure for the visual display of the current situation.
- Tailoring the presentation to the needs of the various echelons: since the SCR's understanding of force is organized into hierarchical levels, the presentation to the users at the various levels can easily be segmented and presented. For example, the JFACC may wish to know only the status of a particular wing, but not the detailed engagement status of each aircraft. The tailoring can readily be done in a rule-based way using the templates (discussed in another section), and the tailoring rules can be modified later as needed.
- Drill down through the SCR levels: the idea of tailoring implies that some beliefs and the evidence to support them will not be displayed even though available in the SCR. The full explanation of a particular belief can be made available by accessing the SCR, that is, by "drilling down" to beliefs (and evidence) at lower levels, even down to the lowest level of supporting sensor data and other data. Since part of the support for a belief may come from beliefs at higher levels, "drill up" will also be made available. The drill up and drill down allow the user at any time to get *explanations* for any item of current belief. This gives cognitive transparency to what otherwise might be a "black box." Of course, like most services of the JBI, this explanation service is subject to the rules of the templates, if the commander chooses to control this service.
- SCR available to all processes: visualization is not the only process that uses the SCR. SCR is the *common* representation for all processes of the JBI, and can be used by any of them; or it can be published to other systems outside the JBI, as the rules allow, transforming data from the SCR to the appropriate native format as needed.

Implementation of the SCR will depend on the middleware, or integrating software, chosen for the JBI. Middleware operation is described in Chapter 6, “The JBI Platform.” Here the functional model for the JBI Input Segment is presented; this is the key method by which the inputs are brought into the JBI and converted to the SCR.

5.4 JBI Input Segment Functional Model

The publication and subscription mechanisms underlying the JBI provide for the dissemination of the information that is published into the JBI. Information flows from a number of sources, and can be input in several different ways. In addition to such inputs, a key feature of the JBI, especially as its evolution continues, will be the ability to use the JBI to reach out for information or capabilities contained in a variety of systems in a manner transparent to the user. Thus, when the user queries the JBI for material that has not yet been published, or the underlying “business logic” of the JBI requires the publication of updated information from existing sources, the JBI will be able to retrieve or otherwise interact with information manipulated by other military systems. This capability will be provided by the use of “information agents” that will allow for flexibility of access and interaction with existing legacy and emerging C⁴ISR systems.²




Figure 13. An evolving “information web” will provide a link from the JBIIS to existing legacy and evolving military C² assets

² The assumption that future military systems will include a number of agents and/or agentized components is backed up by a number of current efforts and reports. For example, a forthcoming report by the Defense Science Board describes the “Integrated Information Infrastructure” for military systems, and proposes that both the implementation and monitoring of this system will be largely agent based. DARPA’s Control of Agent-Based Systems (CoABS) program is directly working toward this vision, and a number of existing military initiatives under way at DARPA and the AFRL focus on agent-based approaches.

5.4.1 Agent-Based Infrastructure

Figure 13 illustrates how an agent-based infrastructure will allow the JBI to evolve toward a seamless integration with existing and evolving systems. Software agents will allow JBI programmers to use fuselets for a wide range of information tasks that include searching, filtering, and collating data from military systems, as well as providing a capability to monitor information in these systems and provide inputs to these systems as authorized by the JBI rules. Code development will largely consist of tailoring these capabilities by instantiating generic capabilities and providing details of what sort of information is required and what system (or generic system type) it might be found in. For example, the code may be as specific as asking for particular entities found in TBMCS, as monitoring information in a network such as the Joint Control and Tracking Network, or querying open sources for a supplier who can provide a piece of equipment needed for logistics support.

To allow these fuselets to provide generic capabilities that are specialized for interacting with systems, they must be based on information about these systems provided in a machine-readable way. This information will be provided using a capabilities language that would provide details of interacting with external sources. These “agent markup languages” are being explored in a wide range of research programs, and DARPA and other organizations are working to make them a reality for the military. In addition, these languages are expected to build on current commercial efforts, such as XML, and to become embedded in commercial tools (as is the current HTML, which is embedded in browsers and page-creation software). These languages are expected to provide mechanisms that will make it possible for agents to easily tailor their interaction with external systems and to be obtainable to the military primarily via COTS acquisitions.




Figure 14. Service-based middleware will allow the interoperability of heterogeneous systems and the retrieval of information from legacy systems

The final key to “reaching out” from the JBI is the development of software tools that allow for the “agentization” of current systems. That is, the systems can be operated by software agents rather than people. This technology, also under development for the military within DARPA and other DoD laboratories, generalizes on previous technologies known as mediators (for databases), wrappers (for web pages and some legacy systems) and facilitators (for agent-based architectures and some sensor-processing applications). This interoperability “grid” will allow systems to share information with each other even where different data or communication standards were used in their designs. It will also allow components running on heterogeneous agent architectures to interact. (It is no more realistic to assume that there will be one agent architecture with a single agent standard than it was to assume that a single data standard could be achieved.)

5.4.2 Service-Based Middleware

Figure 14 shows one approach to building such a grid of agents, under development in DARPA’s Control of Agent-Based Systems (CoABS) program. A piece of middleware, providing a set of interoperability services, is provided for allowing systems to intercommunicate. Systems provide their local information (name, creator, Internet address, etc.) in a configuration file, and the middleware is then able to provide intercommunication with a wide range of other systems. Systems registered on this grid are then able to use the services provided. These services include brokering and yellow pages, which allow systems to find each other based on capabilities; logging and visualization services, which allow for the monitoring, control, and debugging of

multisystem or multi-agent applications; authorization services for security; and a host of others. By providing these capabilities via middleware, the addition of more capabilities (or a reimplementations for greater efficiency or to track an emerging technology) becomes possible without changing each of the systems interacting through the grid.

This combination of agent-based codes, agent markup languages, and an interoperability infrastructure that enhances agent (and legacy) communication provides an “information web” structure that goes beyond the specific needs of the JBI. However, the study team sees this infrastructure as a military necessity, and the study team joins the Defense Science Board and others in endorsing the military development of such an approach. The JBI’s specific ability to interact with a large range of military applications and systems is also greatly enhanced by such an infrastructure. Barring its development, only *ad hoc* point-to-point software solutions, mandating significant code development for interoperability with existing systems, will be able to provide the needs of the JBI.

5.5 Fusion of Incoming Information

It is often said that when users work in modern information spaces, they need software for “creating information from data” and then, once they have the information, “making knowledge from information.” These phrases refer to levels of information fusion. To military fusion specialists, the former is “data fusion” or Level 1 fusion; the latter is called Level 2 fusion, and is sometimes referred to as high-level fusion (HLF) or “situation understanding.” The output of data fusion is a set of “best guesses” of the identity of elementary objects and their behavior (for example, a tank traveling a road at a particular speed). The output of HLF is a linked collection (often a hierarchy) of “best guesses” regarding the aggregation of various types of forces, and the current behavior of those aggregates (for example, a battalion, a division, and a wing of MiG-29s) stored in the SCR.

Fusion is a fundamental service not only for the JBI but for all users of modern information technology. Facing the immense amount of data available at the click of a key or mouse, users seek assistance in discerning the meaning of all the bits and bytes. Most of the time, a user’s limited resource is no longer information but attention. Fusion processes are needed to focus attention on the right things in a timely way.

5.5.1 What Level of Fusion Should the JBI Support?

Fusion technologists use the term “level” to denote a level of capability. This is not to be confused with levels of the hierarchies used in discussing the SCR. Fusion technologists use the term “fusion” where a military specialist might use “fusion and analysis.”

Level 1 fusion refers to sensor (and sometimes multisensor) processing that gives the “perceptual” indication that an object of interest is present. Level 1 provides support to the warfighter’s *perceptual* processes. The analogy is to a person’s eyes and visual system that see objects in the world and recognize what they are.

Level 1 fusion processes usually rely on complex mathematics, statistics, and physics modeling using computer simulation. They are closely coupled to the sensor technology, so that Level 1

fusion is sometimes developed as part and parcel of some new sensor development project. The output of a Level 1 process is the identification (the “recognition”) of an *elementary battlespace object*, and a report of relatively simple behavior of the object. In the design of the JBI, these will be imported to the JBI as streams of inputs or by subscription.

The next step is *aggregation*—how to get from elementary objects to more complex objects, representing more complex things, and ultimately representing situations. The process involves aggregating the evidence for the complex things as well. *Level 2* fusion processes do the aggregation. Level 2 provides support for the warfighter’s *cognitive* processes.

In Level 2 fusion, “*context*” is introduced (context will be defined and discussed later). Qualitative information enters; the dominant processing steps are inferences, not mathematical calculations; military knowledge is used to make the inferences (as opposed to the heavy use of physics knowledge in Level 1 fusion). This is why the usage “fusion and analysis” is justified. The output of Level 2 “fusion and analysis” is *the SCR described above*. Level 2 fusion for the JBI will be performed by JBI software, but it will be informed by portions of other systems’ “common operating pictures” that are inputs to the JBI (for example, GCCS and TBMCS).

Level 3 fusion attempts to infer *enemy intent* (movements, plans, tactics, strategy, etc.). The input to Level 3 fusion is the Level 2 situation understanding—that is, the SCR plus other relevant data and knowledge (for example, enemy doctrine). The JBI concept for Level 3 includes plan-monitoring functions and will be used to monitor the execution of Blue plans as well as Red actions.

Level 4 deals with collection management. It closes a feedback loop that allows the processes of situation understanding and inferring intent to suggest new sensor taskings or new data collecting (from databases) for the purpose of improving the SCR.

5.5.2 Processes for JBI Fusion

Level 1 fusion processes are typically done by heavyweight applications that reside in systems like the Consolidated Acquisition Reporting System. Similarly, Level 4 processes will likely not be carried out within the JBI. Finally, Level 3 processing is beyond the state of the near-term art. Thus, JBI fusion should concentrate on lightweight systems that operate at Level 2. This section describes a processing framework for realizing Level 2 fusion (situation understanding) in the JBI. Two key JBI technologies referenced—publish-and-subscribe and fuselets—are described elsewhere in this report.

Imagine an SCR “in place,” stored in its internal computer representation (which consists of object hierarchies, sets of objects, or any other structure that makes sense in theater combat situations). It has evolved continuously with the operation of the JBI and now represents the current situation understanding.

In every hierarchy of the SCR, every level (itself an object) consists of other objects that represent entities or concepts and include the evidence (or pointers thereto) for believing that such an entity or concept is really present.

Every object and every hierarchical level subscribes to data and information pertinent to it. If anything changes in the information supporting that object; or if any new information alters the “picture” (understanding) of any level, the object or the level receives an input. For example, if the ATO were to change, the Wing level, by subscription, would be notified (perhaps even the Squadron level, if the rules allow them to subscribe to ATO changes). The wing in turn will publish the change, and these will be received (by subscription) by the objects representing squadrons of that wing.

Fuselets are the software intermediaries in all this to-and-fro messaging. This is what a fusion fuselet does:

- Receives inputs (as a subscriber)
- Processes the current state of objects + the inputs, using *rules*, to
 - Publish a rule-based inference, publish a numerical calculation, or publish some data.
 - Do a process that changes something in the SCR. Often, this action will be to link objects at one level with objects at another in a relationship of “evidential support”—evidence of one object is supporting evidence for another object below or above it in the hierarchy.

It is *important* to note that fuselets can link objects to other objects and levels in other hierarchies. For example, Marine platoons may be linked to the forces at the appropriate level of the Army’s hierarchy in the SCR. The separate hierarchies can thereby be “laced together” by links.

The fuselets have been designed to be pertinent to particular object types (such as a squadron) and levels (such as a wing level). If the fuselet is pertinent to a level, such as it might be an *aggregation fuselet* that proposes new things to be represented at its level (for example, if several enemy tanks are identified and are moving together, the fuselet might propose creation of an object representing the tank group).

Rules are normally thought of as having the form *if* (some condition)—*then* (make some inference or do some action). The *if—then* statement does not have to be black-and-white logic (that is, Boolean); it can be probabilistic (that is, Bayesian): *if —then* with probability *p*.

When a fuselet publishes a change or a new inference, the “consumers” are often objects and levels immediately adjacent in the hierarchy. But there can and will sometimes be “remote” subscribers. For example, one subscriber could be a visualization process; another could be a tasking process for getting additional collections of data into the JBI.

5.6 Summary of the JBI Input Segment

This chapter gives a technical description of the JBI Input Segment. Some of the functionality described here may overlap with functions performed by the Interact Segment or the JBI Platform and may ultimately be designed and implemented with one of these other segments.

Nevertheless, the core functionality that must be implemented includes

- Protecting the JBI from spurious or destructive inputs
- Bringing into the JBI volumes of diverse knowledge needed for the tasks at hand using both conventional and agent-based methods
- Transforming information coming into the JBI so it is stored in the JBI's SCR
- Fusing information where feasible to eliminate confusion or duplication

The state of the art in the technical areas supporting the Input Segment is such that a fully featured prototype of the Input Segment could be built today. Improvements will come in spiral development of the JBI as fusion and agent-based technologies improve. The Services should not rush to develop an SCR; this must be coordinated across the spectrum of JBI users, since it to some degree sets in stone the standards for representing JBI knowledge.

The next chapter examines the way the SCR described in this chapter is updated when new inputs or knowledge become available. The JBI Platform is the middleware foundation of the JBI supporting fuselets and the JBI's publish-and-subscribe mechanism.

Chapter 6: The JBI Platform

6.0 Introduction

The previous two chapters described interfaces and inputs to the JBI. In this chapter the foundation of the JBI—that portion providing many of the underlying services—is described. More specifically, this chapter describes how the JBI is constructed and used and the ways the JBI supports the publish, subscribe, query, and control operations depicted in Figure 15. The chapter starts by showing how a new JBI is activated and how it evolves over the course of a mission. A brief overview of the JBI design philosophy leads into a discussion of how the JBI is used to transform information into knowledge useful to a mission. The third major section presents the structure of the JBI from three different perspectives. First is an operational view, which treats such questions as how a new JBI is activated and how it evolves over the course of a mission. The second is a conceptual view of the functions that the JBI Platform provides to clients. This view is intended to show how clients use the platform services to achieve the necessary operational effects. A final view presents some thoughts on the design of the JBI platform itself. These views are not independent; many concepts thread through them all, such as information assurance, access control, bandwidth management, and connections to legacy systems.




Figure 15. The JBI Platform provides infrastructure JBI services: publish, subscribe, query, and control

6.1 Overview

The information staff is responsible for operating a JBI—for standing it up, maintaining it, and modifying it throughout the course of a mission. Because the staff, trained as information

management professionals, serve in several roles, they must embody expertise ranging from system and network administration to the mission-critical information structures embodied in the JBI.

First, some definitions:

- Users: the people who work with the JBI to exploit it
- Client: a computer that uses the JBI as part of a mission
- Server: a computer that provides platform services necessary for the JBI to run
- Platform: the core JBI services that support JBI capabilities
- Process: a computing activity that resides on a machine
- Machine: the computer or communications devices used to support computing or communications functions
- JBI administrators: the information staff who work with the JBI to stand it up and operate it

To simplify its administration, the JBI will contain information about itself—its structure and the structure of the mission. For example, the JBI will record precise definitions of all the information objects it uses, configuration and version information for all of its clients and servers, and an enumeration of the military units participating in the mission and their characteristics.

The JBI enhances information storage and flows among the people and computer processes engaged in conducting a military operation. Improved ability to sift and distill information rapidly provides better guidance to the commander, staff, and warfighters of a mission. A mission is configured with the right information-processing resources, both humans and computers, to manage the mission's information and develop useful knowledge from the information. The JBI's role is to store or provide access to sensor information, intermediate results, and ultimate knowledge in a repository so that it can be shared throughout the mission—subject to proper access authorizations. The JBI also arranges to route information to the right destinations, alerting the people and processes that should respond to new data. The chain of alerts constitutes a workflow process, designed and adapted to process the mission's information. These JBI mechanisms ensure that information is an asset to the mission.

The JBI mechanizes information management using five basic services: publish, subscribe, query, transform, and control. The design of the JBI makes a distinction between the computer programs that make use of these services—"clients" of the JBI—and computer programs that furnish the services—"service providers." The term "JBI Platform" describes the services and the service providers that furnish them to JBI clients. This is a classic client-server design. Clients are largely independent of each other; they are able to share information by storing "information objects" in the JBI that can be accessed by other clients. Service providers, on the other hand, work closely together to create the effect of a giant dynamic library housing JBI objects, accessible to all authorized clients.

The design of the JBI is based on a small set of properties:

- Information is represented in formats that are known to all (or most) JBI clients.
- Each piece of information—or “information object”—is augmented with “metadata tags” that describe the information. For example, temporal-spatial metadata identifies the time and location to which the information applies.
- Information is routed to people and computer processes based on their announced needs. A JBI client “subscribes” to information it needs by identifying information types and ranges of metadata tag values that characterize its needs. From that time on, whenever an information object that matches the subscription is “published,” the JBI routes the object to the subscriber. In this way, “publish and subscribe” mechanisms automate dynamic information flows.
- In addition to subscription-based information routing, each information object is stored in a repository—a digital library—that can be searched by “queries” issued by JBI clients. Complementing subscriptions, these searches are intended to be used by analysts performing knowledge-management functions that cannot be fully automated.

This chapter shows how these properties are embraced by the technical design of the JBI. It shows how information is represented in the JBI, and how JBI clients use and manipulate the information. It goes on to sketch how the JBI Platform services can be implemented. Chapter 7 provides a discussion of the COTS and GOTS components available to build the JBI.

6.2 Transforming Information Into Knowledge

The ability of the JBI to develop higher-level knowledge is due both to automatic processes and to people who analyze and refine information. Automatic fusion programs operate in the JBI by subscribing to the information objects they require as inputs and by publishing fused results. These fusion engines are very powerful but often difficult to modify. To address these shortcomings, the JBI also enables more flexible automatic processes—called *fuselets*—that can be deployed and modified by the information management staff during a mission.

6.2.1 Fuselets

Fuselets are JBI clients that create new knowledge derived from JBI information objects. A fuselet typically subscribes to objects it needs for input, so that as soon as such an object is published in the JBI, the fuselet is triggered. It examines the information in the objects and publishes one or more new objects as a result. These new objects represent knowledge that the fuselet has been able to derive from its inputs. Fuselets can be used to encode a commander’s standing orders, such as “if a ballistic missile launch is detected, issue a ‘major threat’ alert”; “if information arrives about Yeltsin, republish it with keywords ‘Kosovo political’”; “if enemy tanks start to cross the Barada bridge, alert the engineers to blow it up.”

Fuselets are intended to be easy to adapt because of the way in which they are specified. For example, a fuselet can be written as a “script” in a simplified programming language. The script might subscribe to objects that specify the quantity of jet fuel available at each of several air bases and determine whether the total fuel reserves are below a critical threshold that the commander has specified. If so, the fuselet will publish an “alert-for-commander” object

describing the problem. This object may in turn cause other effects—certainly the commander or the commander’s staff will have subscribed to the alert object so that it will show up on their displays.

Fuselets can also be specified using “rule-based systems,” in which simple decision logic can be expressed in a natural way. For example, a rule system might say something like:

- For objects of type airbase-fuel-supply sum object.fuel-on-hand into total
- For objects of type airbase-fuel-supply if object.fuel < 1000 then one-base-low = object
- If total < 3000 or one-base-low then publish object with
 - Type = alert-for-commander
 - Message = “Total fuel at all bases < 3000”
 - If one-base-low then message = message & “Fuel low at” & x.base-name

Although fuselets can be created from scratch using tools for writing and testing scripts and rule sets, most will be extracted from a library of previously built fuselets. These may then be configured or modified slightly to meet the mission requirements.

Fuselets are only as capable as the scripts or rule systems on which they are based. Unless directed by specific rules, they will not exhibit “common sense” or “judgment” when processing information. Although technology may improve over time, military missions cannot depend on fuselets alone for deriving knowledge from JBI information.

6.2.2 Analysts

A human analyst can do the same thing as a fuselet—publish new information objects derived from other information entered in the JBI—but can apply human judgment and problem-solving skills to the task. While fuselets are able to generate new knowledge rapidly, analysts are able to generate new knowledge of far greater depth and “semantic” value.

An analyst will use one or more software tools to interact with JBI objects: to view intelligence reports, situation reports, target lists, messages describing the commander’s intent, overhead imagery, UAV images, maps, and so on. These objects will have been published in the JBI by their corresponding sources. Some of the analyst’s software tools may be specialized for viewing or analyzing particular data types—for example, overlaying enemy position reports on an overhead image with objects of interest identified. The analyst will also have tools for publishing new information objects—perhaps tools as simple as a browser with facilities for authoring documents or forms. The analyst might be able to take information objects that have been analyzed and drag and drop into the new report, to be tagged as sources leading to the derived knowledge. Someone looking at the new report may extract these sources to “drill down” to find raw information behind the analyst’s conclusion.

Analysts will be able to subscribe to information objects using the software tools they use for interacting with the JBI. If a new intelligence report arrives for a region that the analyst is assigned to cover, the tool will issue an alert—perhaps an icon for the new information object will pop up on the screen. The analyst will be able to craft subscriptions to meet information

needs. The tools will also let the analyst issue queries to the JBI—to search the universe of JBI objects to find those that might be relevant to the question that motivated the query.

The analysts thus can function much like a fuselet: responding via subscriptions to new information and publishing knowledge derived from that information. Analysts are not required to use subscriptions—they can browse the JBI, looking for background information for later use or assembling collections of objects to keep handy for analysis tasks, and so on.

The JBI can support a wide range of information transformation processes. Large fusion engines, fuselets, analysts—all are resources for transforming information into mission knowledge. A commander can deploy these resources flexibly to meet mission needs. The structure of “anchor desks,” each staffed by an analyst assigned to a particular role in information management, is easily enabled by the JBI. The ability of the JBI to store information objects and to publish, subscribe, and query them provides a common platform to support arbitrary flows of information in a network of analytic processes.

6.3 Information Structure

The information objects in the JBI are organized into an “information structure” designed to support a military mission. Included in this structure are information objects that support the mission, such as maps, target lists, intelligence background studies, battle plans, and so on. In a sense, the JBI serves as a well-stocked library of knowledge that supports the mission. But the JBI holds another kind of information as well—the information objects that support the execution of the mission’s components. For example, objects will describe force structures, operational units, weapons capabilities, inventories, readiness, and so on. These information objects are designed and installed with an explicit structure and standard operating procedure required to carry out the *operations* of the mission. For example, each unit participating in the mission will be required to publish a readiness report in the form of a specific information object. Fuselets and other automated processes will depend on these reports of individual units in order to create aggregated reports. Likewise, reports of fuel and ammunition supplies will be processed automatically. In a complementary way, participating units will expect to find specific objects placed in the JBI by their commanding units—for example, the ATO. Automated processes within a unit may decipher these orders and determine what actions the unit must take. The JBI will perform its intended function only if all the publishers and subscribers behave according to plan—that is, if subscribers receive the information they expect. In contrast to the “mission library” role of the JBI, this second role might be characterized as the “mission database.”³ Although this dichotomy is not exact—indeed, many objects will serve in both library and database roles—it is a useful explanatory distinction. The JBI platform makes no such distinction: all objects are treated in a uniform way.

³ The JBI is not required to retain *all* of the data or databases that support a mission. Individual units will retain private databases, often housed inside existing C² systems. The information handled by the JBI is the subset that is (1) necessary to be shared with other units or functions within the mission, using the JBI as the sharing mechanism; and (2) couched in standard information object formats that are widely interpretable.

The key elements of the JBI's information structure must be subject to widespread agreement among the joint Services. Standards for information objects and their interrelating structure are similar to today's message format standards (including such things as the ATO) and database standards (for example, the Common Data Environment). This basic information structure will be "built in" to C² software that uses the JBI—that is, acts as JBI clients.

For a given mission, designing, building, and maintaining the information structure of the JBI is the role of the information management staff. This function ensures that the information structure of the JBI is appropriate for the mission. As a mission is stood up, the information management staff must build an appropriate JBI information structure. As units join the mission, their C² systems must be integrated into the mission JBI. As the mission unfolds, if changes to the JBI are required, for example, new information objects must be introduced, or new fuselets deployed, or new versions of C² software must be deployed, it is the information management staff's responsibility.

6.4 Technical Structure

The JBI Platform is the protocols, processes, and common core functions that permit participating applications and organizations to share and exchange critical mission information in a timely manner. It provides uniform rules for publishing new and updated objects into the JBI and promptly alerts any JBI clients that have subscribed to such objects. These properties enable dynamic information flows among client programs of the JBI, serving to integrate the clients to conduct a single mission.

The JBI Platform supports a dynamic digital library of information objects. JBI *clients* are computer programs that may publish objects to the JBI Platform "library" and may be notified when new objects are available. Clients may query for, and subscribe to, objects meeting specified criteria. The client issues a query to identify all objects in the JBI repository that meet the desired criteria. A subscription can be thought of as a "standing query" for specific objects, automatically providing relevant objects to the client when they become available.

The basic element of the JBI is the information object. Technically, an information object is a set of attribute-value pairs, but this representation can be used to express arbitrary data structures. Each information object has associated *metadata* attributes to aid in the query, subscription, and management processes. The JBI Platform has associated with it, in addition to facilitating the access of client programs to the information objects, a set of management processes dealing with their storage, access, and security.

End users access the objects and services of the JBI through the JBI client programs. The client programs implement the mission-important functions of the JBI, such as intelligence management, analysis, integration, and other functions to support the information services needs of the mission. The client programs use the JBI Platform services to get their work done.




Figure 16. Schematic JBI Platform structure

These essential elements of the JBI Platform are illustrated in Figure 16. A client that wishes to subscribe will record with the subscription broker a description of the objects it seeks (signified by a document with a question mark). When a client publishes an information object, the document is entered in the repository and its metadata description is made known to both the query broker and the subscription broker. If the metadata of the new object matches a subscription, the subscription broker will alert the client that entered the subscription. In the query broker, the metadata is saved in an index so that when a client subsequently queries the JBI, the query can be matched against the metadata of all objects stored in the repository. The JBI Platform also provides a fuselet server to manage and execute fuselets created by clients (note that fuselets are technically JBI clients, but the JBI Platform provides a fuselet server as a convenience). Finally, a set of control services is used for monitoring and maintaining the JBI Platform. Included in these are access control services, which ensure that all client accesses adhere to information security policies adopted for the JBI.

6.5 Clients

A JBI client is any computer program that uses the JBI platform services. Although there is a wide variety in the roles and functions of these clients, they all have access to JBI platform services. However, the JBI platform may be configured to allocate different classes of service to

different clients: some may require fast response, some may have low priority, some may not be permitted to access all the objects in the JBI, and so forth.

JBIs clients may use other services besides those offered by the JBI platform. The universal connectivity that the network provides lets these services do such things as access Web resources, read files from computers all over the world, and issue queries against databases. In other words, they can do anything that computer programs can do and, in addition, they have access to the JBI platform services.

Clients take on many forms; and a number of these have been described earlier in this report. Many different programs may choose to contribute to a JBI; some are summarized here:

- *Fuselets*. The notion of “fuselet” is intended to connote a small program that publishes JBI objects by refining or fusing information in a relatively simple way. It subscribes to JBI objects that pertain to the information it is designed to publish. It may also make JBI queries or interrogate other information sources available over the network. Fuselets are part of the “glue” that make the linkages that cause information in the JBI to flow to the right places in the right forms. Some fuselets will be obtained from a library, configured, and placed in service to accomplish a particular job in a JBI. Others will be created using scripting languages or simple programming tools (for example, Visual Basic or JavaScript) to adapt the JBI information flows to needs that arise in the course of a mission.
- *Browsers*. Some users will want to browse information in the JBI, much the way they browse the Web today. The browser will be able to fetch objects from the JBI, to search it using queries, and to make visual presentations of many of the common JBI information objects. A browser may also include tools for creating objects, which are then published to the JBI.
- *Interaction*. Many clients are designed to connect humans to JBI information in special ways. For example, a mission rehearsal client might query the JBI to obtain details of a mission, then build a 3-D fly-through environment in which a pilot can rehearse the mission. A logistics officer might prepare a spreadsheet that uses macros to extract information from JBI objects that record fuel stores at different bases. Or a staff officer might use a presentation tool that can fill in screen templates with data obtained from objects that are returned by JBI queries. An important client might be called an “approval” tool, which presents an object representing an action order to a commander, who then “approves” it; the tool then publishes an object that records the authorization.⁴
- *Legacy C² systems, attached to the JBI with “connectors.”*⁵ When a legacy system develops certain kinds of new information, the connector extracts the information from the database of the legacy system, recasts it as a JBI object, and publishes that object in the JBI. Connectors can also extract information from the JBI and deliver it to the legacy system. Sometimes the connectors are referred to as “wrappers.”
- *Fusion*. One of the missions of the JBI is to encourage fusing information from a variety of sources into high-level “knowledge” that is readily accessible to the commander and other staff. Fusion

⁴ There will need to be a “design pattern” for action objects and how approvals are recorded. Are there two distinct types: unapproved-action and approved-action, or is an approved action merely a copy of the unapproved action with a suitable “approved” attribute and signature?

⁵ This term is chosen to align with the similar concept used in Enterprise Integration Technology products. See Section 7.1.2, “Technologies for Building the JBI: Component/Middleware and Related Core Technologies Roadmap.”

programs designed to take advantage of the JBI can subscribe to information from many different sources: if the information is available, the fusion program will be informed via the subscription, and the program can aggregate the information. In other words, fusion programs can now be written to take advantage of any available information they can understand and evaluate.

- *Sensor feeds.* A sensor may collect information and publish it directly in the JBI. Doing so will allow a variety of fusion programs to find the sensor data and incorporate it into their decision making. The data will also be available for human viewing.⁶
- *JBI-to-JBI gateways.* When the JBI notion is fully deployed, there may be instances of JBIs operating concurrently, each devoted to a different major mission. Almost certainly, some information in one JBI will be needed by someone in another JBI. A gateway can serve as a client of both JBIs, subscribing to information in one and publishing it in the other.⁷
- *Gateways to coalition partners.* Links to coalition C² systems are similar to links to other information sources: a JBI client can extract coalition information and publish it in the JBI; or it can extract JBI information and transfer it to the C² system of a coalition partner. If human authorization is required, this tool can present information to a person and obtain authorization before it is forwarded to the coalition partner.
- *Gateways from other information sources.* Most JBIs will be outfitted with clients whose job is to obtain information from a non-JBI source and publish it in the JBI. For example, United States Message Transfer Format messages could be addressed to a JBI instance. The addressee is simply a JBI client that converts the message into a JBI object and publishes it in the JBI. Another example might be a program that subscribes to a news feed and publishes relevant news stories as JBI objects. This program might have sophisticated means to extract keywords or to generate summaries. Finally, an analyst might be given the responsibility to scan the World Wide Web for information of a certain kind and republish it in the JBI so that it could be found by other JBI clients. The analyst would use a software tool to prepare a JBI object that either copies or refers to the Web resource, then publish the object in the JBI.
- *Proxies for special devices.* Some devices will not have enough computing resources to be full-fledged JBI clients, or may not be accessible using Transmission Control Protocol/Internet Protocol (TCP/IP) network protocols, perhaps because they use special communications links. These devices can use a *proxy client*, running on a computer that is big enough and well enough connected to be a proper JBI client. The proxy is also connected to the device. In effect, the proxy client uses the device as a remote user interface for the proxy program.⁸
- Tools for maintaining the JBI information structure.

⁶ Some sensor data may be very large. The JBI platform makes provision for handling large data sets, usually by insisting that the system creating the data set also serve as the repository for the data; in this way, it need never be copied to another repository. JBI clients access the data directly from the originating repository. Protocols for accessing JBI objects will need to be able to access large objects in portions of manageable size.

⁷ Alternatively, it may be desirable to federate JBI instances within the JBI platform services—that is, to permit JBI clients to obtain information from any JBI that is federated with its “main JBI.”

⁸ This idea is used by today’s pagers to provide e-mail access. The pager cannot support Internet email protocols (SMTP) both because the computer in the pager is too small and because the pager’s radio communication links do not carry TCP/IP traffic. Instead, the pager company operates an e-mail proxy on a computer that is connected to the Internet. When the proxy receives e-mail, it extracts the text message and uses the specialized radio broadcast network to send the message to the pager of the intended recipient.

To the JBI platform, all of these clients are the same—they are provided the same services. In other words, the platform is not aware of the functions that the clients perform, whether they are automatic computer programs, whether people are operating them, or whether they are direct feeds from sensors. Nevertheless, the platform is aware of certain differences between clients—for example, some require fast response, and some are granted greater access privileges.

6.6 Objects

In the JBI, data are recorded and information is made available in the form of information objects, which here are simply called “objects.” These objects are statements about the real world, describing things, events, or plans and intentions. Very often they describe some opinion about reality, instead of reality itself. There may be several objects “about” the same real-world entity, distinguished by “who says so” and “when did he say it.”⁹ JBI objects may be simple or complex. They may contain structured, semi-structured, or unstructured data. One object may refer to others. For example, a JBI object could be any of the following:

- Vehicle position report
- Recorded UAV video
- Commander’s intent
- Friendly order of battle
- A description of available fuel stores

The JBI may contain more than one object “about” the same real-world thing. For example, if two sources report the position of the same enemy ground unit, there will be two objects describing this unit’s position. The objects can be distinguished by publisher and publication time.

JBI objects are put into the JBI when a JBI client publishes them. The JBI platform ensures that the object is stored in a reliable repository. Published objects are available to other clients by name, or through a *query*, which returns a specified collection of objects, or through a *subscription*, which returns a stream of new and changed objects matching a specification.

A JBI object may be modified by its publisher. This is much like republishing a new edition of the same object. Clients that subsequently issue a query will retrieve the new object contents. Clients with a matching subscription will be informed that the object has been modified. A client with a copy of the (old) object will be unaware of the change unless it has used an appropriate subscription to detect changes.

A JBI object may be deleted by its publisher. Subscribers will be notified, and subsequent queries will not return the deleted object.¹⁰

⁹ The JBI can still support a common view of the battlespace. The point here is that it must support *more* than a single view. It must be possible to express inconsistent statements made at different times or by different sources.

¹⁰ Deleted objects may still be accessible depending on policy decisions related to archiving requirements.

6.6.1 Object Model

The JBI object model includes identity and inheritance, but not encapsulation, methods, or polymorphism. In this sense, JBI objects are like XML documents, instead of “objects” in the sense of object-oriented programming or distributed object computing.

6.6.1.1 Object Schema

Every JBI object is an instance of an *object schema*, which in other technology areas might be called a “class,” “object type,” or “data model.” The object schema defines and abstracts a category of things; it says, “For this sort of thing, these particular facts are important and are stored.” That is, the object schema defines a list of *attributes* for which a particular object instance supplies *values*; it is like a template for an object. The object schema also defines the meaning of each attribute and the domain of values. Attributes may be mandatory or optional, and may be arranged in a hierarchy. The XML standard and the forthcoming XML Schema standard illustrate (and may prove to be appropriate for) JBI objects and JBI object schema. However, the important part of JBI objects is the interface they present to clients—the “questions” they can answer—and not the particular syntax used to store and transmit the object contents.

```
<JBI_type> ATO </JBI_type>
<JBI_identifier> tbmcs-99-AX4003 </JBI_identifier>
<version> 0 </version>
<publisher> tbmcs-99-mccarthy </publisher>
<publication_time> 0102001400Z </publication_time>\]
<language> EN </language>
<ato>
  <air_operations_data>
    <day_time> 020200Z </day_time>
    <quantity> 6 </quantity>
    <country> US </country>
    <subject_type> FTR </subject_type>
    <aircraft_type> F16 </aircraft_type>
    <track_number> 401 </track_number>
    . . .
  </air_operations_data>
  . . .
</ato>
```

Figure 17. Example object: an air tasking order using XML encoding for the object, with an embedded XML-MTF encoding of the ATO

JBI object schemas are an essential part of data interoperability between publishers and consumers of JBI objects. An object schema must allow a consumer to *discover* the existence of an object class that may satisfy the consumer’s information need. It must allow the consumer to establish a *semantic match* between the attributes of the object schema and the facts required. That is, it must allow the consumer to determine, “Values of these attributes will satisfy my needs.”¹¹ And it must permit the consumer to cope with any *representation mismatch* between the published data and the consumer’s desired format, preferably by means of automated data

¹¹ Equivalence, or “means the same,” is not necessary here. For example, your data about passenger vehicles may be acceptable for my data needs about all motor vehicles, even though the converse is not true.

mediation. These data interoperability concerns are especially important when the consumer uses the object contents as input for automated processing. In such a case, the object schema is being used as a system interface, bringing most of the problems associated with interfaces, and requiring the same degree of care to ensure correct interoperation.

Object schemas used by the JBI are stored in the JBI itself. In this way, any client can obtain precise and accurate definitions of objects. When a client starts up, it will usually interrogate the stored JBI object schemas to determine whether the JBI definitions of objects it uses are consistent with its own definitions; that is, that there is an adequate match between the client's view of an object and that shared by all other JBI clients. Clients may be able to adapt to some degree of mismatch, as outlined above, but if not, the client may abort. Usually this means that the client software must be upgraded to cope with newer versions of object schemas used in the JBI.

Individual JBI object schemas may be part of a larger integrated data model. This captures relations between objects and allows a client to correctly combine data from different objects. It is tempting to say that there should be one common data model capturing all the relations between all the objects. However, this single-standard approach has been attempted in the DoD data management program and found to be unworkable, so it is not likely to succeed in the JBI. A better approach is to construct and manage a hierarchy of subject-area domains, where each domain represents a community of users and clients who agree on a collection of concept definitions (that is, an ontology) and on a set of schemas defined in these terms.

This approach is part of an overall data management process. If the DoD data management process is changed to become compatible with this approach, then a separate and redundant schema management process for the JBI will not be needed, and connecting legacy systems to the JBI will be much simpler.

Domain ontologies and object schemas are essential when military units describe the information they must obtain from the JBI and the information they will provide to the JBI. Many of these business processes can be worked out in advance. These information exchange requirements are expressed in terms of the types of JBI objects that will be published and subscribed to. Much of the semantic matching between participants can be done and information shortfalls identified in advance. As a result, when a unit "connects" to the JBI, very little manual work will be required to establish the predetermined information flows.

6.6.1.2 Object Metadata

Metadata attributes are used to describe a JBI object itself rather than the real-world entity the JBI object describes. For example, every JBI object has a type, a version number, a publisher, and a publication time. Almost every JBI object includes attributes from a common set: geospatial reference, pedigree or lineage, subject keywords, language, etc. The document properties defined in the Dublin Core are good examples of metadata attributes and possible candidates for JBI metadata.

Metadata attributes are especially useful for describing the objects to be returned by queries and subscriptions, and for efficient evaluation of these object-selection patterns. There will be a common set of JBI metadata attributes for all JBI objects, some mandatory and some optional.

That is, a JBI object does not have to include metadata describing its subject—but if it does, it must use the “subject” metadata attribute. There may also be local extensions of the JBI metadata attributes that define useful search criteria for collections of related objects within the JBI.

Each object must be labeled with a “type” by specifying a value for a mandatory “type” metadata attribute. The type identifies the object schema that describes all the object’s attributes and their meanings. The JBI itself contains a registry of all the object types it uses.

Table 2. Example Metadata Tags

Attribute	Value (* means mandatory)
JBI_type	name of object type*
JBI_identifier	unique identifier string*
version	version number*
publisher	identity of creator of the object*
publication-time	date and time*
keywords	words to use for searching
derived-from	list of object identifiers of antecedents
geospatial-reference	representation of 3D spatial region
temporal-reference	range of times that apply
language	ISO language identifier (EN, FR, ..)

6.6.1.3 Object Encapsulation

JBI objects do not encapsulate or hide data; all of the values for all of the attributes in an object are accessible to any client that obtains the object. JBI objects do not include methods or functions. That is, an object cannot depend on containing executable code for correct interpretation of its values, nor is there any sort of run-time polymorphism.

6.6.1.4 Object Identity

Every JBI object has a unique identifier. Clients can retrieve an object through its identifier. Relations between objects may be established when a value in one object is the identifier of another. Object identifiers may be names; that is, they may be humanly intelligible in whole or part. They do not contain information about physical location, in order to allow the JBI platform to have flexibility in deploying and locating its object repositories. In this sense they are like handles (see <http://www.handle.net>¹²) and not at all like URLs.

6.6.1.5 Object Versioning

The publisher of a JBI object is allowed to change its values. Republishing an object creates a new version of the object that asserts, “This is what I say about the world now.” This does not

¹² The Handle System[®] is a distributed computer system that stores names, or handles, of digital items and which can quickly resolve those names into the information necessary to locate and access the items.

change what was said about the world “then.” The previous version may be retained¹³ and retrieved by JBI clients on demand. Clients may retrieve a specific version number, the latest version, or the version that was current at a specified time.

6.6.1.6 Object Pedigrees

Much of the information in the JBI will be derived from other information. An analyst’s report will be based on intelligence reports, images, and background documents. An aggregate readiness report will be based on reports from individual units. To allow users to “drill down” from an object to find detailed information, objects may contain references to the objects from which they are derived. This “pedigree” information will be recorded as a metadata attribute.

6.7 Platform Services

The JBI Platform is a set of services provided by a distributed structure of computer servers and networks. These services implement the internal operations of the JBI, they provide access for management and control of the JBI operation, and they furnish the services that computer programs use to work with the JBI. This is a conventional client-server design, in which service provision and control are concentrated in the servers, and a wide variety of clients may contact the servers over the network to obtain JBI services.

The JBI services are designed to work together consistently, to be managed and controlled together, and to scale together. It is for this reason that the term *platform* is used here to describe the collection of services. The idea is that a JBI client sees the platform as a consistent set of useful services—all that’s needed to work with a JBI. If more performance or more storage capacity is needed for a JBI, the JBI platform is implemented so that these adjustments can be made easily.

The principal services of the JBI platform available to clients are *publish*, *subscribe*, and *query*—operations that have been introduced earlier. In this section, these operations are described in greater detail, to provide an understanding of how JBI clients may use the platform. However, discussion of how the platform services are implemented is deferred until Section 6.8; indeed, there are a great many implementation possibilities.

6.7.1 JBI Services in Action

This section presents a brief scenario in which a planner collects information relevant to a fighter strike mission in a “cup,” which the planner then passes to the wing, and eventually to the pilot carrying out the mission. The idea is that the cup contains information that the warfighters will need in order to carry out the mission. Moreover, if new information arrives—before or during the sortie—it should be added to the cup. If the information is critical, it should go straight to the cockpit display.

¹³ JBI policy determines whether previous versions are retained, and for how long. Archive storage is the responsibility of the JBI repository.

Here is one way this concept could be implemented, given the conceptual structure of the JBI platform and clients sketched.

1. The planning tool creates a “container” object and places within the container references to any objects the planner has identified—target information, overhead pictures, UAV video, etc. The planner could do this with a tool that uses a drag-and-drop technique to place information objects in the container. The container, when it is published in the JBI, contains references to each of these objects.
2. The cup is now passed to others. A planner at the wing, for example, would refine the mission plan with waypoints or communication frequencies. These refinements are recorded as changes to the cup or to the objects referenced by the cup.¹⁴ As a last step, the wing planner instantiates and configures a fuselet—a “mission fuselet.” The fuselet will subscribe to any new information pertinent to the cup’s geospatial-temporal extent. The intent is that the fuselet will collect new information relating to the mission described in the cup and change the cup accordingly.
3. When new information matches the fuselet’s subscription, the fuselet executes and determines the relevance of the new information. The fuselet might be set up to consult a human to make judgments about some kinds of information. But other kinds, like the discovery of a new surface-to-air missile site near the flight path, the fuselet might arrange to send directly to the cockpit.
4. To reduce the traffic to the cockpit, the wing planner has created a special container with the job of simply containing all information to be forwarded to the cockpit of the fighter that has been given the cup’s mission; this would be a “shooter cup.” On board the fighter is a client that subscribes to the shooter cup. When the mission fuselet wants to make information available in the cockpit, it adds the information to the shooter cup, and it is immediately transmitted to the cockpit.

This design filters new information—automatically or via human intervention—before it is sent to the fighter cockpit. This reduces traffic on already-taxed communications links and provides just the right information to the pilot. There are many possible variations on this theme; for example, the fuselet could prepare a detailed visualization of the new information and transmit this—rather than the information object—to the cockpit.

6.7.2 Publish

The JBI platform provides services for publishing, updating, and deleting JBI objects. A client prepares a JBI object as a set of attribute-value pairs. The client then uses a JBI service to publish the object; this will make the object available to other JBI clients. Later on, the creator of the object may republish the object in order to record changes to it—that is, to change the value of an attribute or to add new attribute-value pairs to the object stored in the JBI. Finally, the creator of the object may invoke a JBI service to delete the object; thereafter, the object will not be available to other JBI clients, but it will have been saved in an archive for auditing or other post-mortem analysis.¹⁵

¹⁴ For this example to work, objects must be able to be changed by clients other than the original publisher of the object.

¹⁵ Some objects may not be archived, depending on policies associated with the JBI.

6.7.3 Query

A JBI client can search the JBI to find objects that match a search pattern; this is called a *query*. The simplest form of query finds the object with a specified object identifier.¹⁶ The result of such a query is to return to the client a copy of all the attribute-value pairs that define the object.

More advanced queries use search patterns that specify values for certain attributes; these are called *attribute patterns*. In principle, the JBI could be designed to search for arbitrary patterns of attributes and values. In practice, however, certain attributes will be chosen as *common searchable attributes*, and the JBI platform will be designed to ensure that searches on these attributes can be executed quickly. Examples of common searchable attributes are the type of the object, the object identifier, and the creator. Because of its mission, a common searchable attribute is the geospatial-temporal region associated with an object. Efficient searches based on such attributes will require that the JBI repositories and search engines use specialized data structures and algorithms, so the design of both clients and services will depend on the way these attribute values are represented. In some sense, the common searchable attributes will be built in to the JBI.¹⁷

A JBI query will return a set of objects that match the attribute pattern specified by the client. In some cases, the query pattern will not have been able to express precisely the set of objects the client seeks. In these cases, the client will need to process each of the objects returned by the query to determine whether the object is one that the client wants. The JBI query services are intended to do the bulk of the work in searching so that clients will typically have only a small job remaining.

6.7.4 Subscribe

A JBI client can subscribe to find out when an object is created, updated, or deleted. The client prepares an attribute pattern to describe the class of objects it wants to learn about. Whenever an object published in the JBI matches the subscription pattern, the subscription is said to “trigger.” When the subscription triggers, it informs the subscribing client about the object that triggered the match.

A subscription returns the attribute-value pairs of the object that triggers the subscription. To control bandwidth, the client may specify that only a subset of the attribute-value pairs be reported. For very large objects, such as images, the platform may provide services that allow clients to retrieve only parts of the object so as not to overwhelm communication channels with data that a client ultimately will not use. In some cases, a client may not require any of the attributes of the object that triggered the subscription; it may simply want to be activated when the subscription triggers and will obtain data from other sources or by issuing a query to the JBI.

¹⁶ The query uses the identifier as a search pattern. It may also request a specific version number; if no version number is specified, the latest version is returned.

¹⁷ Technically, it is the types of the *values*, not the attributes, that may require built-in features in the JBI. JBI query services may also want to offer specialized searching for attributes that have arbitrary text values: searching for words in the text, or Boolean patterns of words, perhaps with the ability to rank-order the results.

6.7.5 Transform

The JBI platform includes services for running fuselets and other critical JBI clients that participate in the JBI workflow processes. These fuselets have the same status as any other JBI client, but they are housed within the platform principally so that they can be monitored and controlled by the information management staff.

6.7.6 Control

In addition to the services provided to JBI clients, the JBI platform implements services that are used to monitor and control the JBI itself. JBI clients will use some of these services, but the information management staff, whose job is to maintain the JBI, will use most of them.

6.7.6.1 Access Control

The military nature of the information in the JBI demands strict ability to control access to data. The access control services are used to define who may retrieve or subscribe to which data and JBI object types and at which levels of classification. To verify the identity of JBI clients and users, the JBI must *authenticate* them. Suitable *authorizations* specify the services and data that clients and users may use. The JBI will have tools that allow the information management staff or security personnel to register or expel users and to change authorizations.

It is important to realize that JBI clients as well as users must be authenticated, because it is essential to establish the authenticity of all information in the JBI. An object that claims to be data from “the radar on Mt. Auburn” is believable only if the sensor system that published the object can be shown to be the JBI client it claims to be.

All of these issues fall under the broad categorization of “information assurance.” Clearly, the success of the JBI will depend on a sound system for information assurance. Unfortunately, this is a difficult and longstanding problem of considerable importance for the DoD and for which no new insights are offered here. As access to information improves, access control becomes more essential, and DoD programs in this area must be supported.

6.7.6.2 Configuring, Monitoring, Managing

The information management staff is responsible for configuring, monitoring, and managing the JBI using control functions implemented in the JBI platform together with interactive applications that allow administrators to visualize and modify system properties. The spectrum of these activities is described above in Section 4.3.

Keeping the JBI running smoothly is considerably harder than today’s “network administration” task. Two perspectives illustrate the difference:

- The JBI platform services are delivered by collections of computers and networks working together. Ensuring that a service is working properly requires more than simply ensuring that each of the computers and networks is working. For example, several servers, at different geographical locations and linked by a communications network, will usually provide a JBI’s object repository. If the repository needs more capacity, not only must a new server be added to the network, but it must then be introduced into the repository federation and begin to share the load. If a communication link between two of the repository sites has insufficient bandwidth to carry the

traffic required to coordinate the operation of the repository, the problem must be tackled as one of tuning the repository service. An alternative to obtaining more communication bandwidth might be to reconfigure the repository servers.

- The operation of the JBI depends on information flowing due to dynamic publish-and-subscribe linkages. If some aspect of this automatic flow is not working properly, the information management staff will need to diagnose and fix the problem. Perhaps a fuselet has failed and is no longer subscribing to and processing its inputs. To monitor and repair the JBI, the information management staff will need to be familiar with the information structure represented in the JBI and will need tools to inspect objects, clients, and fuselets that participate in the structure.

The JBI platform software must be designed to operate in a 24×7 mission-critical environment. In addition to using techniques to provide robustness in the presence of equipment or software failures, the software will need to monitor its own operation and report “alarms” to JBI information management staff.

6.7.6.3 Bandwidth Management

The JBI depends heavily on a communications infrastructure of limited capacity. Demand on the communication networks is likely to vary widely, depending on sensor tasking, crisis conditions, damage to communication links, and so on. Ensuring that the JBI continues to service its mission role requires managing the full spectrum of JBI resources, including available bandwidth, to ensure that the quality of service required for each information flow is adequate. For example, it might mean dynamically rerouting lower-priority information traffic in order to carefully control the use of slow links to airborne platforms to ensure that mission-critical traffic gets through.

The JBI is designed to relieve some of the more obvious communication bottlenecks. For example, the object repositories will stage copies of heavily used objects on servers close to the clients that use the objects. This avoids duplicate transmissions, especially over critical long-distance links.¹⁸

6.7.6.4 Object Type Management

The JBI contains a repository of *type definitions*, including object schemas. During the course of a mission, it may be necessary to define new object types and ensure that these definitions do not conflict with others. Tools for object type management are available when the information management staff must interrogate the type repository or develop new object types.

6.8 Platform Implementation Possibilities

Providing the JBI services in a way that meets the mission requirements implies a number of needs, both functional and nonfunctional. Functional needs were discussed above, and involve support for client programs to identify, access, manage, and modify information objects in a dynamic environment.

¹⁸ This kind of staging and link management is illustrated by Wide Area Assured Transport Service.

Nonfunctional requirements include

- *Reliability and robustness*: critical information for a mission must be available when needed, in the presence of failures and attack.
- *Scaling*: the JBI includes data such as sensor outputs, implying a very large number of information objects.
- *High performance*: use of the JBI to support dynamic C^2 implies the need for rapid access and dissemination of data.
- *Intelligent bandwidth utilization*: critical information must be disseminated even when available communications is limited.
- *Measurability and manageability*: the commander must be aware of the status of the JBI. Problems need to be quickly identified and repaired.

The key idea in implementing the JBI Platform is to create the effect of a single set of JBI services while distributing the implementation over multiple servers. Thus the illustration in Figure 16 is schematic; in reality, the elements shown in the figure may be replicated, partitioned, and distributed. The distributed structure addresses a number of the requirements mentioned above (such as scaling and robustness).

An important strategy in the distributed implementation is *load partitioning*. If all requests for JBI services were to be funneled through one server, or even through a small number of servers, the system would not perform adequately. Small objects that must be processed quickly to participate in near-real time mission-execution loops would be delayed by the processing of large objects that have less stringent deadlines. To address this problem, the JBI assigns clients to servers dynamically so as to meet performance requirements. When a client registers with the JBI, it provides information about the types of data it will publish. Based on this information, the identity of the client, and policies established by the information management staff, the client is told which servers to use when publishing which objects. Likewise, the client is told where subscriptions should be recorded. In this way, separate servers can handle the fast, small traffic. Moreover, if the information management staff needs to adjust assignments, clients will adjust their behavior. These dynamic configuration processes are largely invisible to a programmer writing client software: the programmer invokes a publish or subscribe function in a “JBI software library,” which deals with the intricacies of load partitioning.

The distributed implementation of the JBI services also permits efficient information management without requiring clients to become involved. For example, patterns of subscriptions can be used to guess which clients will use a newly published object and to stage a copy of the object near the clients—or such staging can be controlled via policies set by the information management staff.

6.8.1 Technology Antecedents

Implementation of the JBI builds on capabilities that have already been demonstrated by commercial systems and research projects that are well along in development. This section reviews the relevant technologies of digital libraries, enterprise integration technologies, and

XML. The JBI implementation will depend on other technologies that are already widely deployed by the military, such as networking connectivity (TCP/IP) and services, client-server techniques, and the Web.

The JBI can be built incrementally alongside existing services. It is not intended to replace existing services, but to supplement or integrate them. Moreover, implementation of the JBI can exploit systems already in place. For example, the Web and Web browsers should serve as one access path into JBI information. A portal server can provide a “Web interface” to JBI information, translating JBI information objects into types known to Web browsers. To further exploit existing Web systems, the JBI should probably allow many of its data type standards (for example, HTML, JPEG, and PDF) to be used in JBI information objects.

6.8.1.1 Digital Libraries

The JBI is a digital library with some special characteristics. A research program organized by DARPA and the National Science Foundation has been developing digital library technologies for several years. One definition of a digital library is “the collection of services and the collection of information objects that support users in dealing with information objects, and the organization and presentation of those objects available directly or indirectly via electronic or digital means.”

Digital libraries provide these services by federation¹⁹—a design in which different elements achieve concerted action by coordination using network communications. Each object in the library is uniquely identified by a name, or *handle*. A library will have one or more *repositories* responsible for storing digital objects—usually documents—and for enforcing any access rights the author requires. Associated with each object is some *metadata* to specify properties such as author, title, creation time, keywords, and document type. One or more *index servers* are responsible for indexing documents, allowing a query based on metadata or (in some cases) document contents to return handles for the stored documents. Finally, a *handle service* figures out, from a document’s handle, which repository holds it. All of these services—repositories, indices, and handle managers—can be replicated and federated to form a large, distributed library.

Many of the functional capabilities for the JBI may be found in digital library and enterprise integration technologies. However, these require enhancements to satisfy the JBI needs, both nonfunctional and functional.

- *Digital object definition.* Understanding the nature of the objects in the JBI, their relationships, and how they can be disseminated is important to having the “virtual repository.” The digital library community has been working on this issue and developed, for example, concepts that distinguish the streams of audio from the digital objects that contain them and the disseminators that allow them to be “viewed.”
- *Location-transparent naming.* If the JBI is to support a wide variety of applications, providing them with the information needed for their individual functions, there needs to be a common way of naming objects. Furthermore, because storage locations will change (due to failures as well as

¹⁹ Barry Leiner, “The NCSTRL Approach to Open Architecture for the Confederated Digital Library,” *D-Lib Magazine*, Dec. 1998, <http://www.dlib.org/dlib/december98/leiner/12leiner.html>.

normal reconfigurations), it is important that such naming be independent of the actual storage location and that there be a way of resolving names into locations and other properties (for example, access control or object typing) for retrieval at that time. The Handle System is an example of a location-transparent naming system that can provide this function.

- *Repository Access.* Both the digital library and the JBI have the concept that the “user view” is independent of the storage of the objects. This is facilitated by having standardized methods for accessing (depositing, retrieving, modifying) objects in the various storage sites. The Corporation for National Research Initiatives, for example, has developed the concept of the Repository Access Protocol to address this issue.
- *Information Organization.* One of the main characteristics distinguishing a digital library from the Web is the “organization and management service” that a library provides. Similarly, the JBI must not be simply a random collection of information. It must be organized and structured to support the mission. Concepts used for digital library organization can be applied to the JBI.
- *Metadata.* The digital library community has done extensive work to define metadata standards and approaches.²⁰ Metadata will be critical for the JBI to support multiple applications and organizations in a known way.
- *Search and retrieval.* Unlike the World Wide Web, which tends to support search based on simple keywords, the digital library community has developed search techniques that exploit the information organization and metadata. It also has developed concepts like index services and collection services that facilitate fast and efficient searching. These will be very helpful in developing a JBI that can support the military requirements for intelligence and planning.
- *Integration of legacy databases and systems.* Digital libraries include the ability to interrogate databases (converting, on the fly, database contents into digital objects). Standard object-oriented techniques (for example, Stanford’s STARTS protocol) will prove useful as the JBI evolves and attempts to integrate existing systems.

The digital library technologies do not meet all the needs of the JBI. Some aspects that will need attention are

- *Object types.* Only a small number of fixed types of documents are held in a digital library. There is no provision for dynamically introducing new object types and managing the schemas, ontologies, and other information required to interpret objects of the new types.
- *Number of objects published and rate of publishing.* The JBI is expected to produce a great many objects at a high rate. Although digital libraries can scale to become large, they are not engineered for real-time publishing.
- *Linking.* Although digital library objects refer to one another, the links are sparse and not particularly vital. By contrast, JBI objects will make heavy use of object references, and they must be right. This means, for example, that repositories will have to provide transaction control for updating collections of objects consistently.
- *Indexing.* Digital library indices are intended principally for metadata with textual values. The JBI will require, in addition, ways to index geospatial-temporal values, and perhaps other specialized domains.

²⁰ Stuart Weibel, “The State of the Dublin Core Metadata Initiative, April 1999,” *D-Lib Magazine*, April 1999, <http://www.dlib.org/dlib/april99/04weibel.html>.

6.8.1.2 Enterprise Integration Technologies

“Enterprise integration technologies” refers to a variety of middleware techniques for integrating otherwise separate enterprise (business) information technology applications. Although these may take many forms, generally they involve software “connectors” to provide a linkage between each application and a shared structure for communicating information from one application to another. One form of linkage is to route messages from one application to some or all of the other applications; this form is called message-oriented middleware.²¹ As an example of a message-oriented approach in a business setting, an order-entry application might send a message announcing the details of each order to a billing application and to one of two fulfillment applications, depending on which product is being ordered.

One particular form of message-oriented middleware uses a publish-and-subscribe metaphor. One or more applications may publish an *event*—a short message containing information from the publisher—that is communicated to all applications that subscribe to events of that type. The middleware makes sure that events are delivered to all subscribers.

Although the publish-and-subscribe model is central to the JBI, existing middleware does not precisely map to the needs of the JBI. Two of the differences are

- The JBI applies the publish and subscribe terminology to objects, whereas industry’s middleware applies the terms to events; these are quite different. Merely delivering an object’s attribute-value pairs as an “event” will not work.
- Industry’s “subscribe” does not use matching based on attribute patterns—in fact, it does no matching at all. Instead, publish and subscribe operations are applied to named *channels*: for example, the “order channel” is assumed to carry events signaling new order entries. Using existing channel mechanisms to achieve the effect of matching is awkward and would lead to inadequate performance.

The notion of “fuselet” is similar to “component software” schemes being advanced as a way to simplify enterprise application development. Perhaps the best known of these is Enterprise JavaBeans, which sets out rules for writing software modules (beans) that can be assembled into larger applications. This technology defines an “application server” in which beans may be executed, together with a suite of middleware services (messaging, transacted database access, publish and subscribe) available to the beans.

6.8.1.3 XML Technologies

The World Wide Web has created demand for enterprise application integration “in the large”—linking electronic commerce and other applications on behalf of businesses all around the world, not just applications under the control of a single information technology organization. The explosion of the Web has also overtaken some of the technologies that launched it, such as the HTML document format. These two needs have converged in a technology called XML, which provides a uniform way to describe data structures in a textual form. Data structures expressed in XML may be used to represent documents (hence an HTML successor) as well as structured data (useful for conveying order entries or other business data). These uses have given rise to XML

²¹ Another form uses connectors to link each application to a shared database.

variants for a wide range of purposes; for example, the real estate industry has developed Real Estate Listing Markup Language, a way to represent and exchange multiple listings of property for sale.

The XML thrust is useful to a JBI implementation in two ways:

- XML can be used to represent JBI object schemas and objects. Because of the huge commercial interest, COTS XML tools will be readily available. Already vendors are showing browsers that can present XML data in a wide variety of different ways, useful for JBI data visualization.
- As part of the XML effort, groups will convene to define data standards. DoD may wish to contribute to these standards, with a view to exploiting the COTS tools that ensue. For example, it might be especially worthwhile to work with metadata standards, so that digital libraries and indexing schemes might evolve to suit military needs. An important special case is the geospatial temporal metadata: there are both military and commercial needs for such tagging. By cooperating on metadata definitions, it might be possible to stimulate the digital library, commercial imagery, and Geographic Information System communities to produce commercial systems of considerable value to the military.

XML is likely to have a huge impact on military C² systems integration independent of the JBI. Already, for example, XML is being used to encode the standard military message formats, including the ATO.

6.8.1.4 Military Developments

Many of the critical ideas in the JBI have been demonstrated in military research prototypes or deployments. This section mentions only a few that relate closely to the JBI technical structure.

- *Information Dissemination Management*. This program shows the power of the publish-and-subscribe model attached to a large digital document repository. A large number of documents, images, maps, and other objects are catalogued according to standardized metadata formats. Users may enter subscriptions so as to be notified when new documents are catalogued. The implementation is a distributed system intended to deal with large numbers of large documents; it stages objects at sites according to subscriptions, and it includes infrastructure for bandwidth and communications management (Wide Area Assured Transport Service). This system and its predecessors have found enthusiastic use in the Balkans engagements. Although it demonstrates many of the essential elements of the JBI, it has two shortfalls: (1) the publish and subscribe linkages are not very fast; and (2) it is not designed to support application integration.
- *Visage*.²² This system, developed by Maya Designs as part of DARPA's AutoBrief program, has developed an object representation called *U-forms* that is applicable to the JBI. Objects are attribute-value pairs with no methods; objects are stored in a repository available to clients; all clients can parse all objects; clients can subscribe to objects in order to sense changes. Visage uses this infrastructure to share data among collaborators. It illustrates both the suitability of the object representation and the way that publish-and-subscribe can carry information to the people who need it.

²² <http://www.maya.com/visage>.

- *XML-MTF* (AC²ISRC/C²P).²³ The power of a uniform data representation is illustrated by this project, which encodes all the military Message Text Formats 101 into XML documents. Not only can these documents be easily parsed by COTS XML parsers, but they can be readily viewed in different ways by prototype XML document viewers. XML is likely to have a rapid and widespread impact as a uniform way to represent military information.

6.8.2 Challenges and Enhancements

There are challenges in applying the technologies from digital libraries and other emerging areas to the JBI. These primarily derive from the dynamic nature of the JBI and its environment. Digital libraries are relatively static. Objects are created at a fairly slow rate (a human rate) and inserted into the various repositories and registered through a process that, of necessity, needs to be optimized for speed of resolution but not necessarily for speed of administration.

In the JBI, on the other hand, large numbers of objects (for example, target tracks) are created rapidly and need to be made available to the users rapidly, in near-real time if not real time. Developing the extensions to the infrastructure technologies associated with digital libraries (for example, the naming system) to support the rapid administration of the objects requires research.

Another dimension of the dynamics has to do with failure modes and communications. In the digital library environment, if a document is not available temporarily because of a server or communications failure, that is relatively acceptable. However, that situation is not acceptable at all in a military mission environment. Hence, many of the approaches to storage, infrastructure, etc., would need to be examined and approaches for robustness developed that are matched to the JBI environment.

Publish and subscribe mechanisms imply a degree of dynamics and are not typically supported in a digital library environment. How they could be integrated with the digital library infrastructure requires investigation, although these are more straightforward than the two issues above.

6.9 Essential Elements of the JBI

This chapter sketches a concrete design for the JBI in order to provide a basis for debate about its operation and benefits. Many design details are omitted, and it is likely that an implementation may make choices that differ from the ones sketched. The study team believes, however, that the benefits of the JBI derive from the following few essential features:

- A considerable degree of information standardization, in the form of standard object types, is required to allow information to be widely disseminated and understood. These standards must have far greater scope than today's pairwise agreements that allow two C² systems to communicate. New efforts in the commercial sector to tackle inter-enterprise integration, using emerging technologies such as XML, may help DoD.
- Tagging is essential as a means to filter the enormous amount of information produced on the battlefield. Deploying tagging based on a common set of metadata tag definitions will have enormous benefits, even if the rest of the JBI is never implemented.

²³ <http://www.herbb.hanscom.af.mil/download.asp?rfp=R35&FileName=XMLMTF.ppt>.

- “Publish and subscribe” is a powerful mechanism—and metaphor—for routing information. It subsumes point-to-point messaging, broadcasting, and other structures. It enables a rich set of workflow structures to be built, yet adapts easily to new requirements.

It is important to remember that the JBI is not intended to *replace* C^2 systems, but to be the substrate for *integrating* C^2 systems. Major C^2 ISR systems and their operators use the JBI to integrate the information structure of a mission.

(This Page Intentionally Blank)

Chapter 7: Technologies for Building the JBI

7.0 Introduction

The goal of this chapter is to present a comprehensive strategy for constructing the JBI. The investment portfolio spans basic and applied R&D in defense systems and related technologies, vigorous and iterative operational experimentation and prototyping of advanced capabilities, and acquisition and deployment of the JBI to the field.

There is no question that building a system of the complexity of the JBI requires a considerable design and implementation effort. Nevertheless, it not need be built from scratch. There is a significant base of rapidly evolving commercial software and hardware technologies that can and should be leveraged. This chapter proposes an iterative, spiral-oriented development approach to allow rigorous experimentation with new technologies and operational capabilities to drive the refinement of JBI functionality over time. The focus is on defining Defense's "business logic" to support Defense-specific applications built on a strong base of commercial technologies and systems.

While there is much that can be acquired from the commercial sector as building blocks, it is important to recognize that DoD's needs for integrated fusion, planning, and execution systems are not what is driving the development of commercial products and capabilities. This chapter identifies the gap between commercially available technology—today and in the near future—and the JBI technical architecture as described in this document. This chapter reviews Defense's current R&D portfolio and recommends an investment strategy to ensure that the right technological capabilities will be available when the JBI needs them.

The organization of the chapter is as follows. Section 7.1 reviews the current and near-term state of the art in the constituent underlying information technologies and Defense-relevant applications. These are presented as roadmaps covering (1) commercial technologies for components and middleware processing; (2) networking and communications equipment and services; (3) interaction technologies for information capture, presentation, and collaboration; (4) technology development for Defense applications for fusion, planning, and execution systems; and (5) planned developments and deployments of defense C² systems. Section 7.1.2 compares the emerging commercial middleware technical architecture with the JBI technical architecture developed in this report to reveal the shortfall in technologies needed to realize the JBI. The study team's summary recommendations for building the JBI and investing in its constituent technologies are found in Section 7.3.

7.1 Commercial Roadmaps

In this subsection, roadmaps are reviewed for commercial technologies for components and middleware processing (Section 7.1.2); networking and communications equipment and services (Section 7.1.3); interaction technologies for information capture, presentation, and collaboration (Section 7.1.4); technology development for Defense applications for fusion, planning, and execution systems (Section 7.1.5); and system development for Defense applications (Section 7.1.6).

7.1.1 Overview

The JBI has many of the same attributes as, and is not necessarily more complex than, the kinds of mission-critical enterprise systems being developed and deployed by leading commercial enterprises like Federal Express, Wal-Mart, and those being developed to support the Stock Market. The feasibility of constructing the JBI rapidly and cost-effectively depends critically on understanding what is and will soon be available from the commercial sector. It is equally important to understand where commercial technology has a different emphasis and offers different capabilities than what is needed for the JBI.

The study team believes that the JBI can exploit considerable commercial technologies for component-based object-oriented systems, middleware services for applications integration and data storage or retrieval, networking and communications systems and services, and capture, presentation, and collaboration technologies. These are covered in Sections 7.2.2 through 7.2.4.

In addition, DoD has been making important research investments in these areas to further develop and expand the base of constituent technologies for the long term. Many of these investments will bear fruit in terms of commercial capabilities that can be integrated into future iterations of the JBI. This is covered in Section 7.2.5. Additional research investments are being made in terms of prototyping critical Defense applications, such as logistics, planning, and C². These represent opportunities for insertions into the JBI development spiral and are described in Section 7.2.6.

Finally, it is important to understand Defense's current plans for deploying the next generation of operational C² systems, as well as experimentation exercises such as EFX. This roadmap is captured in Section 7.2.7.

Given the study team's emphasis on exploiting commercial technologies, this presents two primary challenges.

First, the underlying commercial technologies evolve at a rapid rate. It is dangerous to select the "best available" commercial technologies too early, as this runs the risk of locking the system into obsolete technology before it is even deployed. Furthermore, commercial technology is not a panacea. Even mission-critical commercial software can be fragile and unreliable. The driving focus on commercially oriented enterprise applications, like customer care and supplier value chain integration, may be a mismatch for Defense requirements for security, reliability, and real-time performance.

Second, the existing methods of Defense acquisition are inadequate for procuring modern systems centered on information technology. What is needed is not a detailed architectural-level specification, but rather a clear statement of desired functionality and the benchmarks by which performance can be measured. This *performance-driven* approach is in contrast to one that produces a detailed technical architecture—often frozen too early—that forms the basis of a complete implementation. This is supported by an agile *spiral approach* that follows an iterative, feedback-based strategy that refines specifications through a process of design, prototyping, and testing.

7.1.1.1 The Challenge of Exploiting Rapidly Evolving Commercial Technology

The development of information technology over the past 50 years has been truly astonishing. There is an often-quoted statement attributed to Thomas Watson, Sr., in the early 1950s that the United States needed about half a dozen computers. Today, many tens of millions of personal computers are sold each year. If anything, the pace of innovation is accelerating.

These advancements appear to be increasing at an exponentially increasing rate. This is codified in Moore's Law, which states that the number of transistors that can cost-effectively be integrated on a single integrated circuit chip doubles every 18 to 24 months. This law has held since the early 1970s, and has been extrapolated through the first decade of the 21st century.

While this growth is true at the component level, it does not quite operate like this at the systems level. Figure 18 captures the evolution of system capabilities as a function of time. Capabilities advance exponentially during times immediately following the introduction of new system architectures. These come about because the rapid advancements at the component level have made possible a completely new and more cost-effective way to organize information technology systems.

Figure 18 captures three breakthrough system architectures: mainframe-centered systems (for example, batch processing, centralized time sharing, and relational databases), workstation-centered systems (for example, personal computing, spreadsheets, word processing, client-server processing, local-area distributed processing, and engineering design applications), and Internet-centered systems (for example, client-proxy-server processing, wide-area distributed processing, and electronic commerce). When first introduced, innovation expands rapidly, leading to the rapid development of new capabilities within systems. But then systems undergo a period of consolidation, with a slow improvement of capabilities.

Eventually the underlying component technologies—processing, memory, storage, and network bandwidth—that are continuing to advance at exponential rates enable a new cost-for-performance breakpoint, and a new architectural alternative emerges. With this new architecture comes a host of new applications that could not have even been dreamed of in the preceding generation.




Figure 18. Information technology evolution and the dangers of too rigid specifications

7.1.2 Component/Middleware and Related Core Technologies Roadmap

Middleware refers to software that glues together, or mediates among, multiple programs. The JBI must be the “glue” for a variety of existing and future C³I systems, so a portion of the JBI platform must consist of middleware. A wide range of commercially available middleware should enable construction of the JBI from a mature COTS base. Unlike some other technology areas, where significant breakthroughs are both required and anticipated, the commercial middleware market seeks instead to achieve “turn of the crank” improvements in time required to successfully deploy solutions, cost of deployment, and scalability and performance.

Component technology refers to software that is packaged in an object-oriented approach. This leads to the opportunity for increased productivity, through facilitated reuse, along with increased performance, quality, functionality, and time to market. Components encapsulate functionality, with well-defined interfaces, and are amenable to the assembly of flexible systems. They are marketable entities: self contained, with introspection and support for visual composition tools, which fosters reuse. Standard components (such as JavaBeans) are interoperable across languages, tools, operating systems, and networks. A variety of commercial tools are available to build and deploy components.

7.1.2.1 Emerging Commercial Technical Architecture

Within the next 5 years, the study team expects current distributed computing environments to rebase on distributed component technology, such as JavaBeans and Common Object Request Broker Architecture (CORBA) components (and to a lesser extent, because of their homogeneous-only orientation, ActiveX components) tightly integrated with the Internet.

An architectural view of commercially available middleware is shown in Figure 19.




Figure 19. Architecture of commercially available middleware

Four backplanes form an essential basis of this architecture:

1. *Web infrastructure.* Middleware supporting the Web, including browsers, viewers, and a variety of extensions to HTML such as XML; enhancements will continue to be introduced rapidly.
2. *Systems management and monitoring.* Multiple vendors provide the capability to remotely monitor, control, administer, and visualize the performance of heterogeneous and geographically disparate systems.
3. *Security.* COTS middleware lacks some of the completeness and maturity of military-level security software. Nonetheless, significant point capabilities are available today, especially for intrusion detection, virus monitoring, and firewall protection. Availability of policy-based security systems and certificate (public key) issuance and revocation systems is currently limited to a small number of vendors; additional policy-based integrating offerings within the next 3 years are anticipated.
4. *Foundation infrastructure.* This set of infrastructure components, including networking hardware and systems, and relational databases, is widely available.

The remainder of the architectural components are described below.

Enterprise JavaBeans and Java Standards: Java standards are particularly useful because they enable relatively easy interaction among a rich variety of middleware vendors, simplifying acquisition and speeding deployment. Look, for example, at the Java 2 Enterprise Edition platform standard. It ensures that the standard enterprise Java technology services automatically work together, enabling COTS messaging services products that use the Java Messaging Service standard, electronic-commerce components based on Enterprise JavaBeans, and JavaServer Pages technologies, to work together without additional programming.

Connectors—leveraging legacy applications and data: Existing legacy software both accepts information and commands and emits information through application programming interfaces. Rather than modify the legacy software, connectors map the existing application programming interfaces into an appropriate object (such as an Enterprise JavaBean or an information object). This approach enables the full suite of JBI software to interact with legacy software as though it were also built with the JBI model. In effect, the connector acts as a “wrapper” around the legacy application, so all data passing through the wrapper, both into and out of the application, are transformed into the appropriate format.

Note that a critical aspect of successful wrapping of legacy software is the design of schema—a standardized view of key data elements and processing functions, and semantic mappings so that an output X of one legacy system corresponds to input X of another (for example, coordinates in GPS or longitude-and-latitude format). The emergence of XML as an industry standard for providing a structured approach to establishing a schema between legacy applications is a powerful development that shows great promise for integrating DoD legacy systems.

Messaging—events, data, and transactions: Message-oriented middleware provides asynchronous message queuing for application-to-application communication, enabling ongoing processing while the message bus delivers the data (as messages). Participating applications can send, receive, and process messages with guaranteed message delivery, even if process, system, or network failures occur.

An *event* is anything of business significance, such as changing a customer address or accepting an order. Events may be processed in near-real time. They are defined at a semantic level that is relevant to the business processes being modeled (in other words, they are more than a data record and typically are described at a higher level than a message). They are decoupled and asynchronous: producers and consumers of events are anonymous to each. Consumers “tune” to a common channel rather than to a specific producer of information.

Producers—which can be individuals, applications, or connectors—*publish* business events on channels. Consumers *subscribe* to channels if they are authorized for access to the channel. Channels are secure and are assigned quality-of-service attributes such as reliable, guaranteed, or transactional semantics.

Business processes are a set of steps or activities required to accomplish a business objective. Examples are hiring an employee, upgrading a credit rating, acquiring a customer, or processing an order. The high-level goal of enterprise integration is the seamless execution of business processes across disparate applications and partners.

A critical aspect of the JBI is to define the information producers and consumers, the events, channels, and business processes relevant to the Defense enterprises being supported by the JBI.

Business rules define how information objects are manipulated and in turn transferred to the next processing step. They are typically specified in terms of events, conditions, and an action to be taken if the specified event meets the specified action. For example, the cumulative expenditure of petroleum, oil, and lubricants might trigger an action that causes a refueling operation to be scheduled. Developing and testing workflow and event modeling software may require

simulation systems to allow the visualization of process steps. Business analysts model processes and use these models to enable the integration of the underlying applications.

The digital library: Digital libraries provide content-independent data storage, permitting heterogeneous objects within a scalable archive. Objects are categorized in a manner appropriate for each object type. For example, a book might be categorized by title, author, or subject; a movie by producer, director, or theme. Parametric searches address metadata entries like author, subject, title, length and the like.

Natural-language queries allow users to express searches in simple, natural style, without concern for exact word positioning or Boolean constructs, and typically return a list ranked by relevance. This requires strong textual analysis to distinguish, for example, between “the White House” and “the white house.” Content-based search capabilities are also available from some vendors, including query by image content to allow searches based on color percentages, position, distribution, and image texture.

A degree of content authentication is available in digital libraries. Referred to as rights management, electronic signatures or watermarks can be encoded onto films, images, photos, and manuscripts to indicate content origin. More important, in the construct of digital libraries, metadata files are essential so that the pedigree of information can be maintained and so that users can quickly ascertain the contents of a library without downloading the entire content. Metadata files are orders of magnitude smaller than the files themselves. NITF-2.0 is an example of an industry and DoD imagery metadata standard for specifying the content and source of digital imagery.

Digital libraries may encompass a variety of database servers: object servers for the digitized content files such as video clips, more conventional databases for catalog information and pointers to the objects. It is possible to construct a commercial digital library so that frequently used objects are stored near the end users independent of the lookup database location, providing higher performance for very large objects.

XML: XML is a data format for structured document interchange on the Web. It is also used for more general exchange of structured data. XML is not a predefined markup language: it is a metalanguage—a language for describing other languages—that allows design of a customized markup. (A predefined markup language such as HTML defines a way to describe information in one specific class of documents; XML allows definition of customized markup languages for different classes of documents.) XML can do this because it is written in Standard Generalized Markup Language (SGML, ISO 8879), the international standard for defining descriptions of the structure and content of different types of electronic document.

XML removes two constraints restraining Web developments: dependence on a single, inflexible document type (HTML) and the complexity of full SGML, the syntax of which allows many powerful but hard-to-program options. XML simplifies the options in SGML and allows development of user-defined document types. A Document Type Definition (DTD) is a file (or set of files), written in XML, containing a formal definition of a particular type of document. It defines the names for element types, where they may occur, and how they fit together. This

capability to define document types tailored to specific communities is key to the applicability of XML to Defense.

The Resource Description Framework (RDF) is a foundation for processing metadata; it provides interoperability between applications that exchange machine-understandable information on the Web. RDF emphasizes facilities to enable automated processing of Web resources. RDF metadata can be used in many application areas. Examples include resource discovery for improved search engine capabilities; cataloging for describing the content (and the content relationships) available at a website, web page, or digital library; intelligent software agents for knowledge sharing and exchange; content rating; describing collections of pages representing a single logical “document”; and describing intellectual property rights of web pages. RDF with digital signatures will be key to building the “web of trust” for electronic commerce, collaboration, and other applications.

RDF defines a mechanism for describing resources that make no assumptions about a particular application domain and do not define the semantics of any application domain. The definition of the mechanism should be domain neutral, yet the mechanism should be suitable for describing information about any domain.

A number of commercial vendors are preparing XML software tools. These fall into three general categories: (1) parsers that can interpret XML and present the document appropriately to the user; (2) tools for creating and managing DTDs that allow user communities to create and disseminate DTDs for specific communities; and (3) editors that support creation of XML documents.

Object databases: Three database technologies are broadly available commercially: flat files, relational database management systems (RDBMS), and object-oriented database management systems (ODBMS). Object databases may be particularly well suited to nontabular data. Commercial object databases typically support Java and C++ objects. Representing such data in an RDBMS may require unique mapping code that is avoided in an ODBMS. In addition, for some data models, relationships may be stored with the object to avoid relational joins, leading to higher ODBMS performance. Scaling and performance of commercial ODBMS solutions are not yet at the maturity level of the much more established RDBMS systems.

Lightweight directory access protocol (LDAP): LDAP was designed at the University of Michigan to adapt the complex X.500 enterprise directory system to the Internet. A directory server runs on a host computer on the Internet, and various client programs that understand the protocol can log into the server and look up entries. Typically, all browsers have LDAP clients built in. LDAP can be used to look up services and devices and to access information across the Internet or intranets.

Portals—web application servers: Web clients are widely available for a variety of platforms, including emerging device types such as handheld digital assistants. These clients gain access to the JBI via a portal, which may include security and personalization support. Portal middleware includes a web application server. Additional software runs with the web application server—for

example, personalization middleware (which interacts with policy and security software to determine which icons to place on the user's browser desktop).

Java standards provide advantages at the portal web server. JavaServer Pages simplify the integration of dynamic content from Java applications within standard HTML Web. JavaServer Pages combines standard HTML tags with JavaBeans components to cleanly separate dynamic content presentation from its generation.

Information dissemination and event notification: Agents and filters define specific searches or extractions on data sets. Filters are generally simple rules, such as “show no images” or “limit geographic range to these coordinates.” Filters are generally synonymous with scripting agents, using strings of keywords united by Boolean logic, and are available today. Intelligent agent technology is not yet commercially mature. Another term that is sometimes used for this concept is *standing queries*.

Once information has been gathered and filtered, it may be pushed or pulled. *Push* indicates that the information is delivered immediately by the producer (for example, via e-mail); *pull* indicates a more passive model, in which the consumer seeks out the data when it wants it (for example, from a web page).

Use-adjustable filters and preferences: Besides the end user-selectable aspects of the information-dissemination and event-notification software, end users may use a simple level of filtering and personalization interface appropriate for a command officer who is not a computer specialist. Examples include indicating key words for searching and check boxes for the type of output formatting (for example, sort by date or relevance score). More complicated and personalized filters and preferences can be envisioned that prioritize types of data displayed, method of display (for example, pie vs. bar charts), and time- or mission-dependent priorities. Furthermore, conditional priorities, such as “show the Predator image only if a cloud-free line of sight is available” or “show JSTARS data within 20 kilometers of my current location” can be established using the technology described above.

7.1.2.2 Component and Middleware Technology Roadmap

The study team's view of the emerging commercial availability of component technology is shown in Figure 20.




Figure 20. *Emerging Commercial Availability of Component Technology*

Current availability: Today, there is widespread commercial availability of tools and middleware supporting component models (for example, Java and CORBA) that interoperate well with the Internet. These tools provide for visual composition of components and web application server support of servlets and Java Server Pages. Servlets enhance an HTTP server by enabling request and response services: when a client sends a request to the web server, the server can send the request information to a servlet to construct a response for the client. Servlets can be loaded automatically when the web server is started, or can be loaded the first time a client requests their services. After loading, a servlet continues to run, waiting for additional client requests. This approach to server-side Java allows sharing of code and applets between server and client applications, makes maintenance of state information easier on server systems, provides a convenient database interface via database access routines, enables access to other applications on the server, and provides access to the Java class library. A rich set of capabilities is provided via implementations of the Java 2 Enterprise Edition platform standard.

Mature availability by 2001: The next 2 years will bring commercial maturity to technologies already under development. Object technology will become increasingly prevalent in real-time embedded systems. For example, the Java 2 Platform, Micro Edition will address the consumer

space, from smart cards to pagers to the set-top boxes. Java embedded server devices with a very small footprint will allow the remote installation and management of new software and services on embedded devices, over the network, dynamically, securely, and just in time.

Mobile support will grow quickly via intercommunicating objects and tools to enable device manufacturers to easily leverage the interoperability benefits of component support. For example, the Jini™ is targeted to make adding an electronic device to a network as easy as plugging in the base unit of a new cordless phone. Jini will allow three key capabilities: *spontaneous networking*, or the ability to dynamically (without drivers) establish communication, sharing, and exchange of services between any hardware or software on a network; *federation*, or the ability to marshal a dynamic distributed system of devices and software components, such as services on phones, TVs, cameras, and computers; and *discovery*, or the protocol for how a new service becomes a part of the federation and advertises its services to other members.

Component-centric middleware will become ubiquitous, providing object implementations to wrap legacy data and application functions, and to bridge between noncomponent and component-based technologies. Enterprise JavaBeans will become the norm for a scalable, distributed, cross-platform component model for reusable business logic.

An electronic marketplace will emerge for reusable components. There will be a variety of publicly available components; an economic model of payment for function maturity, stability, and scalability, will drive commerce in these objects.

Mature availability by 2004: The next 5 years will bring entry-level commercial maturity to technologies that today are research projects. Infrastructures will become available, supporting new uses of interacting components. For pervasive devices (that is, items with a small footprint and low processing power, such as web browser-enabled cell phones, enhanced pagers, and personal digital assistants), this infrastructure will enable object-to-object communication across a wired or wireless distributed network of interacting devices, providing significant new application opportunities.

Adaptive components will monitor performance metrics and will self-tune to meet the real-time needs of the larger system. Similarly, quality-of-service constraints will be designed into the architecture of the object models, so that components can make processing decisions based on policy guidance from policy management directories. The first step will be the availability of component instrumentation for simulation testing to understand performance characteristics under a loaded system.

A “virtual world” model of natural interaction for designing, simulating, and assembling systems will emerge. This notion of replacing manipulation of atoms with electrons will enable more rapid deployment of new hardware systems due to reduced physical build time and simulations.

The electronic marketplace of reusable components will evolve to domain-specific horizontal and vertical component packages, tailored to specific industries. This will affect the nature of application and systems development, because the starting points for building new systems will be more affordable and robust. One of these new systems will be a “hunter-gatherer” system to

visualize and integrate information, leveraging a distributed system of sensors engaged in synchronous and asynchronous communication to provide diverse data-mining and information-fusing capabilities.

7.1.3 Network Technology Roadmap

7.1.3.1 Current and Emerging Network Architecture and Services

The JBI is not viable without a robust and pervasive communications substrate. Although many Defense systems will employ special signaling methods and protocols optimized to their intended purpose, the *lingua franca* within the JBI for communicating data—or in some cases pointers to data flowing in special protocols—will be the IP suite.

IP networks used by the JBI fall into two categories:

1. Purpose-built networks at all scales (system-area, local-area, and wide-area) implemented specifically for the JBI
2. Segments of the commodity Internet, operated both commercially and by government

Thus it is useful to construct roadmaps for commercial Internet hardware and software and for raw communication bandwidth. It is important to understand these trends, as they bound what can be built and what can be bought.

Internet hardware and software and the features they offer evolve continuously and rapidly, primarily to meet the demands of the enterprise, service provider, and (more recently) consumer markets. Many of these developments are irrelevant, or at best peripheral, to the JBI: advances in xDSL (digital subscriber loop) technology, for example, or router enhancements to support IP telephony.

Critical technology trends and developments that directly affect the JBI include the following:

Internet protocol security: Rapid advances in communication technology have accentuated the need for security in the Internet. The IP Security Protocol Working Group is developing mechanisms to protect client protocols of IP. A security protocol in the network layer will be developed to provide cryptographic security services that will flexibly support combinations of authentication, integrity, access control, and confidentiality.

Raw routing speed: The speed of an Internet router is the composition of the bit rate at which its interfaces can drive the attached communication links, plus its internal processing power that ultimately limits the rate at which IP packets can be processed. Routers capable of processing a sequence of the shortest possible IP packets (43 bytes) and delivering them to their attached communication links at rated speed are said to operate “at line rate.” Routers are available today that operate at line rates of 2.4 Gb/s (OC-48). Within 1 to 3 years, line rate operation at 9.6 Gb/s (OC-192) will be commercially available. This means that commercial products for very high-speed wide-area networking will be pervasively available.

IP multicast protocol support: The delivery of a message from a single sender to multiple recipients is a common requirement. If it is achieved by generating at the source as many replicates of the message as there are intended recipients, needless congestion of the

communication links may result; for example, the link leading from the source to the first router in the network carries all the replicates. On a tree of communication links leading from the source to all the recipients, IP multicast technology sends only a single copy of the message and replicates only when the tree branches. Commercial applications of IP multicast include video or news distribution and collaboration or distance learning applications. Commercially available Internet routers support IP multicast, and both higher education and some Internet service providers (ISPs) have begun to use the technology. Within the next few years, the applications of multicasting should proliferate as more ISPs enable multicasting within their networks.

Mobile IP protocol support for mobile network nodes: Internet hosts connected by satellite, radio, or other wireless links need seamless connectivity even when they travel out of range of their point of attachment to the network, also known as a *base station*. Mobile IP works because the mobile node is able to discover whether it is at home or away from home. A host determines whether it is on its home network by using extensions to Internet Control Message Protocol Router Discovery Protocol (RFC 1256). Routers acting as home agents or foreign agents advertise their existence. Home agents are routers located on the mobile node's home network that are capable of tunneling the mobile node's datagrams to it while it is away. Foreign agents are devices on a network that are capable of acting as a detunneling point for datagrams to the mobile node. RFCs 2002–2006 are additional standards documents for mobile IP. Commercially available Internet routers currently implement mobile IP, and the burgeoning wireless data market, driven by the deluge of Internet appliances such as the PalmPilot, will hasten its deployment.

Quality-of-service support: The classical Internet “best effort” single grade of service has proven inadequate for the needs of emerging new Internet applications such as video delivery and IP telephony, unless the network is (usually uneconomically) overprovisioned. Fortunately, the IP protocol suite provides for a type-of-service byte in the IP header, and this byte has been used in a variety of informal and proprietary signaling schemes by ISPs for many years. Examples include giving priority to network management traffic or giving interactive traffic priority over less time-critical services such as e-mail and file transfer. The increased use of data networks for latency-sensitive packet audio and video such as Real Audio or Real Video, as well as the emergence of support for virtual private networks with service level agreements, is driving the rapid development of these capabilities within IP networks.

Resource reservation protocols (RSVPs) and Differentiated Services: The Internet research community has been developing formal standards for providing grades, or qualities, of service. In the first iteration, the Resource Reservation Protocol (RSVP) was developed. It was envisioned that Internet hosts requiring a certain level of service for a session would use RSVP to request and if possible reserve the resources needed for that service in the chain of routers leading back to the correspondent host. It was soon realized, however, that, even at current usage levels, routers in the core of the Internet would not be able to maintain—let alone honor—the reservations for even a small fraction of the millions of sessions passing through the core at any given instant.

Internet researchers have agreed on a modified architecture in which RSVP is used at the edges of the commodity Internet but sessions with identical or similar requirements are aggregated and treated in one of a few possible ways in the Internet core. This architecture is known as “Differentiated Services.” Another approach, applicable to Internet segments built on an asynchronous transfer mode infrastructure, maps RSVP requests onto the appropriate asynchronous transfer mode class of service. Support for both approaches is available at the beta test level from major Internet hardware vendors, and it is expected that Differentiated Services will be widely available in IP networks within 3 years.

Network resource managers and other network-oriented middleware: The ability of Internet hosts in an organization—or, more properly, applications running on those hosts—to request enhanced services from the network clearly poses administrative as well as technical problems: what is to prevent *everybody* from asking? Just how many requests can be honored, and who will be the favored ones? The technical approach to solving these problems posits the existence of an administrative system known variously as a bandwidth broker or policy manager that accepts policy statements from appropriate organizational authority and both administers policy and adjudicates resource requests. Software implementing these functions and communicating with routers is becoming available from vendors, along with application software that is “RSVP-aware.”

The bandwidth broker is but one example of a class of services—also known as middleware—which performs administrative functions that interface applications to the Differentiated Services network. Standards are lacking in this nascent field.²⁴ Nevertheless, study and development are vigorous, and usable software will become available in the next 1 to 3 years, driven by the deployment of Differentiated Services. This is expected to become widespread in the enterprise market and to a lesser degree among ISPs.

Precedence: Under Differentiated Services, the IP header’s type-of-service byte has been renamed the “DS” byte. In the standard²⁵ the bit pattern xxx00000 (where “x” is either “0” or “1”) is used to signal a “legacy” interpretation of the first three bits—the precedence bits. Although precedence is unlikely to be used in civilian applications, router procurements for the JBI or other military applications will be able to specify precedence together with a set of router queuing disciplines to achieve classical precedence results.

Source-based routing: Traditional Internet routing protocols route packets according to their destination address only, even though the source address is also present in the IP packet header. In the increasingly heterogeneous Internet in which some—typically higher-performance—segments are intended only for the use of special communities, this behavior is inadequate. At a typical research network interchange, for example, traffic entering the exchange from a *commercial* subscriber and destined to a university must exit the interchange—and be routed over—the commodity Internet. Entering *academic* traffic destined for the same university destination, however, has permission to be routed over a much higher-performance segment of

²⁴ Cf. draft-aiken-middleware-reqndef-00.txt.

²⁵ Cf. RFCs 2474 and 2475.

the Internet. Source-based routing has been available from router vendors for several years, but generally at the expense of diminished router throughput; some router vendors currently advertise line-rate performance, and it should be generally available within 1 to 2 years.

Communication bandwidth: Current commercial practice over optical fiber uses bit rates of 2.5 Gb/s per wavelength, with 10 Gb/s becoming available. Dense wavelength-division multiplexing technology permits on the order of 100 wavelengths to be sent on an individual fiber today. A 40-Gb/s bit rate has been demonstrated experimentally, and 1,000 wavelengths per fiber is forecast for 5 or more years out.

For geostationary satellites, the NASA Advanced Communications Technology Satellite operates at data rates as high as 622 Mb/s. Current commercial practice does not exceed 155 Mb/s. For satellites in low Earth orbit (LEO), the troubled Iridium network offers a few Kb/s. Teledesic is planning an initial digital data rate of 2 Mb/s, growing to 150 Mb/s—but Teledesic is not scheduled to be operational until 2002, if then.

7.1.3.2 Network Architecture and Services Roadmap

Table 3 shows network capabilities in the present, near term, and far term.

Table 3. Network Capabilities

1999	2002	2005+
2.4 Gb/s line rate	9.6 Gb/s line rate	38.4 Gb/s line rate
2.5 Gb/s fiber bandwidth	100 Gb/s fiber bandwidth	Tb/s fiber bandwidth
GEOSAT: 155 Mb/s LEOSAT: 9.6 Kb/s	GEOSAT: 622 Mb/s LEOSAT: 2 Mb/s	GEOSAT: Tb/s LEOSAT: 150 Mb/s
Multicast and mobile IP router support	Multicast applications and mobile IP proliferate	Multicast applications and mobile IP ubiquitous
Quality of service/RSVP/ Differentiated Services/ precedence router support	Wide-scale deployments in enterprise networks; reservation-aware applications ubiquitous	Wide-scale deployments in ISP networks; reservation-aware applications ubiquitous
Policy-based routing support limited availability	Policy-based routing support widely available	Policy-based routing ubiquitous

7.1.4 Interaction Technologies

In broad terms, interaction technologies may be partitioned into three areas: *capture*, *presentation*, and *collaboration*. *Capture* refers to technologies that support the user in extracting information for the JBI and manipulating it to meet current needs. Examples include barcode readers, wearable computers with voice input, or tools that automatically extract data from television broadcasts. As the name implies, *presentation* covers technologies used to deliver information to users. Collaboration technologies provide a means for linking users, work groups, and teams, including linkage with machine agents resident in the JBI. Based on configuration and integration, some technologies can have a role in each of these functional areas. This section

presents a projected 10-year evolution of interaction technologies to support the JBI. These projections are based on commercial and government R&D.

7.1.4.1 The Current State (1999)

Current user interfaces for information systems are workstation oriented and rely mainly on manual capture technologies such as a keyboard or mouse as a data-entry device. Standard techniques such as drag and drop, cut and paste, form filling, menu item selection, and manual text creation are used to interact with data. E-mail, production tools, and other collaboration aspects of the interface are separate applications. Limited organizational, business process, and user modeling is available to support automation of processing details associated with information and collaboration management and first-order intelligent interface concepts. Moving information from one application to another to support knowledge creation and information dissemination often depends of considerable manual input by the JBI user(s). In spite of these limitations, there is adequate interaction technology to support effective information visualization, rapid formation of virtual work groups, and tools for analysis and decision support. This will be achieved through the use of standard window-based graphical user interface methods. The commander would normally see a map of the AOR, with a separate window showing high-resolution detail for a current area of interest. Typically it will include overlays to show locations and states of critical assets and evolving activities. Other panes may provide decision aids, show a view of the electronic collaborative environment, and represent live feeds from different information sources. As requests for information and analysis are made, the JBI builds fuselets to aggregate information for delivery to the user.

Separate tools, such as e-mail, telephone, and an electronically mediated collaborative environment support negotiation, coordination, and other problem-solving activities through the JBI. These interfaces also support document sharing and interactive (turn taking) electronic sketching (for example, electronic whiteboards) as a means of developing or exploring concepts or coordinating actions.

7.1.4.2 The Near Term (2004)

It is anticipated that speech interaction systems will become the dominant method of interaction with the JBI within 5 years. This will allow the user to interact more naturally with the information environment and, perhaps more important, facilitate the speed of information requests and other interactions. Combined with active templates and semantic integration of ontologies within the JBI, the input system will automatically subscribe to relevant information sources based on the user's request and the system's understanding the operational context and state. Furthermore, it will disseminate reasons for the query and information or analysis outcomes to relevant users at each echelon in the organization, thereby increasing awareness of both the reasons for and understanding of delivered information. Each information packet can be tailored in viewpoint, perspective, and level of detail, based on JBI templates or another suitable internal mechanization. It is also anticipated that mixed initiative collaboration concepts will mature to improve both the speed and quality of planning and decision making. Three-dimensional audio presentation will allow focused speech intelligibility to be improved while

simultaneously allowing a set of users to benefit from opportunistic pickup of situational information from background conversations.

Automation technology should be advanced to the point where it supports intelligent updating and modifying of active templates, user models, and organizational models used for information dissemination, coordination, and information presentation. Individual interaction devices (for example, wearable computers, palm pads, and gesture recognition) are expected to increase in availability and functionality. This will provide continued connectivity as users move to different job sites, support the devices' use in field activities, and provide a means for the JBI to capture, in real time, critical information derived from the field activities. Because many of these devices should become unencumbering by 2004, the interface will be more transparent.

Continued improvement is expected in the ability to integrate audio, video, and text data sources into fused objects for presentation. It should be possible to exploit expected improvements in bandwidth management technology interactively with active templates to determine effective information presentation for a distributed set of JBI users from command centers to field sites.

7.1.4.3 The Far Term (2009)

It is reasonable to believe that in 10 years the JBI interaction technologies will be able to intelligently track and form a fairly sophisticated “understanding” of a complex campaign as it unfolds. As a result, interaction with the JBI will be more flexible and adaptive. These gains are contingent on achieving gains in user, task, and organizational modeling technologies.

Mixed-initiative interface concepts should be more robust, and useful forms of adaptive interface technology should be available to help manage user workload and improve situation awareness and problem solving. Sophisticated cross-language transformation methods should be available to support a more seamless collaboration among users and with JBI information sources and tools in a coalition environment.

It is also expected that the JBI will be able to exploit near-real time formation of dynamic simulations to aid visualization, analysis, and decision making. (It will also support mission rehearsal.) While the study team expects some form of simulation capability to be fielded prior to the 10-year point, the dynamic construction of a simulation derived from multiple, distributed sources is expected to take several years to mature and probably will not be available until 2009.

7.1.4.4 Summary

This section recaps the anticipated development of interaction technologies.

7.1.4.4.1 The Current State (1999)

The technology is

- Workstation oriented
- Windows-based graphical user interface for information capture, presentation, and collaboration
- Crude intelligent aiding techniques and methods
- Intrusive portable devices for field site connectivity, information capture, presentation, and collaborative support

7.1.4.4.2 The Near Term (2004)

The technology will be characterized by

- Speech-based interaction with single-source 3-D source localization as the main mode of interaction
- Mixed-initiative interaction technology as a method of intelligent aiding
- Integrated multimedia presentation
- Unencumbered portable devices, including wearable computer technology for mobile or field application

7.1.4.4.3 The Far Term (2009)

The technology will be characterized by

- Real-time composable simulation technology for visualization
- Robust user, task, and organizational models to support mixed-initiative and other forms of aiding
- Multisource 3-D sound location technology

7.1.5 Technology Development for Defense Applications

In addition to COTS, a variety of existing GOTS and research-related software projects are potential sources of systems components for the JBI. This section provides a representative sampling of relevant projects to show that most of the technologies on which the JBI relies are being developed. Although many existing projects should continue and even receive enhanced support, some projects may appear duplicative to widely available COTS products. As candidate technologies are reviewed during JBI development, it may be found appropriate to combine projects or to eliminate some altogether to avoid duplication of effort.

7.1.5.1 DARPA

7.1.5.1.1 Information Technology Office

Active Networks: This program develops network architectures that are responsive to application requirements and ultimately programmable by applications traffic. The intent is to enable networks that are more dynamic and adaptable to the needs of applications.

Tolerant Networking Technologies: This program develops networking technology to enhance the network's ability to heal itself in response to service outages. The intended result is network systems that are more resilient to network failures.

Broadband Information Technology: This program develops technology for application to nonmilitary fiber-optic networks to render them more usable for military applications. Examples include very high-speed cryptographic hardware that can keep pace with the high data rates of these systems.

Next Generation Internet: This program develops the key networking technologies for the generation after next—for example, at data rates beyond 10 Gb/s. The program is based on a series of testbeds and planned demonstrations in collaboration with academic and industrial participants.

Communicator: The program is working on text-to-speech synthesis to speak machine output, speech recognition to hear and understand spoken dialog, and voice recognition to verify speaker identity. The goal is to develop wireless, mobile dialog interaction that provides context tracking without a keyboard.

Intelligent Collaboration and Visualization: The goal of this program is to develop technology to support planning and execution. Key components are collaboration among distributed systems connected with diverse bandwidths and accessed through a range of devices from handheld to room sized; collaboration among persons with sporadic connectivity and with changing personnel; and technologies to select, sort, and search a multimedia environment.

Information Survivability: The goal of this program is to develop technologies that can be used to create survivable systems. The desired system will detect malicious and suspicious activity and can be used to guarantee minimum essential continued operation in the face of attacks.

High Confidence Computing Systems: This develops modular, verifiable prototype systems and technologies to provide integrated, flexible, and adaptive support for the security, dependability, and real-time requirements of distributed mission-critical applications.

High Confidence Networking: The security of defense computer communications systems depends on having sophisticated assurance methods available at all times. This research area develops the protocols and management software that will guarantee that all the critical communication features can be used effectively even when severe stress is placed on the network.

Survivability of Large Scale Systems: This subprogram develops technologies for credible detection of intrusions and suspicious events, to allow infrastructure elements to detect events, to allow damaged systems to redirect to the most important tasks, and to allow compromised systems to reconfigure to a state not susceptible to the original attack.

Wrappers and Composition: This subprogram develops tools to allow easy insertion of barriers to attack into legacy systems and to enable assessment of the security and survivability of such strengthened systems.

Quorum: Defense applications require seamless interoperability, distribution over multiple nodes, and the sharing of information in support of rapidly organized joint and coalition missions. Such systems will be composed of rapidly evolving COTS assets of diverse capabilities deployed in hostile environments. Today's commercial technology emphasizes functional interoperability with virtually no control over performance. Trends are exacerbating the problem as the commercial world has adopted a model of "distributed" computing based on least-common-denominator desktop systems, and no single system provider supplies technologies that span the full range of infrastructure to meet Defense needs.

Sensor Information Technology: The program will create the binding between the physical world and cyberspace. Today's information systems focus on human input or computer-generated data for fodder, but the future will build on continuous streams of real-world physical data. The program is founded on the concept of a networked system of cheap, pervasive platforms that

combine multiple sensor types, reprogrammable general-purpose processors, and wireless communication. The result is as though a supercomputer were miniaturized and distributed into the environment, with each node computing and collaborating to “see” into its sensor region.

Information Management: The information management program seeks long-term fundamental breakthroughs in our ability to acquire, organize, use, and preserve valuable information. It strives for major advances in acquiring and effectively using vertically integrated as well as horizontally distributed information resources to provide the defense analyst with a comprehensive ability to assess a rapidly changing situation.

Intelligent Collaboration and Visualization: This program will exploit commercial information technology products such as Java, HTML, and VRML to develop, demonstrate, and deliver advanced collaboration and visualization technology to the military through inclusion in commercial products and through insertion into the military information systems development pipeline.

Translingual Information Detection: The program’s mission is to develop the technology to enable individuals working in English to locate, access, and use network-accessible text documents in multiple languages without requiring any knowledge of the target languages. This will require advances in component technologies of information retrieval, machine translation, document understanding, information extraction, and summarization, as well as the integration technologies that fuse these components into an end-to-end capability yielding substantially more value than the serial staging of the component functions. The mission extends to the rapid development of retrieval and translation capability for a new language of interest.

7.1.5.1.2 Information Systems Office (ISO)

Active Templates: The goal of this program is to develop a scalable, simple, distributed software infrastructure for mission planning and execution by integrating symbolic problem solvers and external data sources with simple but expansive visual interfaces. Active Templates will result in an innovative integrated mission management technology to enhance military planning and execution by reducing planning time, speeding recovery of problems encountered during execution, and increasing quality of operational control for battlefield commanders.

Advanced Logistics Project: This project focuses on developing and demonstrating advanced information technologies that will allow the Air Force to get control of the logistics pipeline and the entire logistics business process. It will define, develop, and demonstrate fundamental enabling technologies that will allow logistics and transportation assets to be deployed, tracked, refurbished, and redeployed more efficiently than ever before.

Information Assurance and Survivability: The program will develop security and survivability solutions for the Next Generation Information Infrastructure that will reduce vulnerability and allow increased interoperability and functionality. Technologies are now being developed in the areas of prevention, detection and response, and security management. Ultimately, these will be integrated into a security architecture that, while integrating security and survivability concepts, techniques, and mechanisms, will also provide interfaces for future security upgrades. The

program has just recently begun a \$200 million, 4-year initiative to improve the availability and security of the JBI. The program comprises four subprograms: dynamic coalitions, science and engineering tools, autonomic information assurance, and cyber command and control. A brief program description follows.

The *Dynamic Coalitions* program seeks to develop technologies that enable secure collaboration within dynamically established mission-specific coalitions while minimizing potential threats from compromised partners and external attackers. Topic areas: multidimensional security policy management, secure group management, and coalition infrastructure services.

The *Information Assurance Science and Engineering Tools* program, by creating a science-based design and assessment environment, will allow both DoD and commercial developers to create systems with understood assurance properties and measurable effectiveness against attack. Topic areas: cyberscience, information assurance metrics, mathematics and models, science-based methods for information assurance design and assessment, and integrated environment for information assurance design and assessment.

By creating an autonomic defensive control system, the *Autonomic Information Assurance* program will allow systems to encode tactics to handle known attacks and high-speed attacks so that the decision makers can focus on the innovative and sophisticated strategic situation. Topic areas: light autonomic defenses, modeling, integration, experimentation, response selection techniques, assurance posture specification, policy projection, response mechanisms, and system state estimation.

To enable leaders to understand the situation and act quickly to thwart strategic information warfare campaigns, the *Cyber Command and Control* program seeks to create a cyber defense decision support system. Topic areas include

- *Situation awareness*—Derive and represent information describing the state of the system in terms of actual and potential attacks and their impact on the information and operational functions supported by the system.
- *Course of action development and execution*—Determine and evaluate possible coordinated responses to the current and projected attack situation, using available defensive mechanisms and possible system functional realignments. Carry out a selected course of action, invoking or providing directives to appropriate defensive mechanisms and responding conditionally as contingencies are experienced.
- *Forensics and damage assessment*—Analyze what information, system functions, and decisions have been affected by adversary activities. When a new or unknown type of attack is suspected, examine system state, software, and supporting data, to determine its means of operation and to identify ways of removing, isolating, or otherwise countering it.
- *Integration*—Develop common CC² architecture within which above capabilities can function. Lead efforts to rationalize and fulfill common information needs and other shared aspects of the research projects. Plan and organize coordinated experimentation.
- Together, these programs should form the basis for significant improvements in assuring the availability of the JBI information infrastructure and Security of the information residing there—even in the face of determined adversaries.

Rapid Knowledge Formation: This will produce the advanced technology required that enables end users to directly enter knowledge to create massive knowledge bases (1M axioms) in less than 1 year. This technology is essential if knowledge bases are to capture the expertise of human subject matter experts and allow that expertise to be applied rapidly, in real-time C^2 systems, to complex problems that have never before been encountered. End-user knowledge entry requires tools that enable AI novices to directly grasp what is in a knowledge base and compose formal theories without having to understand formal logic. In addition, this goal implies a requirement for parallel entry of knowledge by teams of 25 to 50 individuals. The knowledge bases to be created will be demonstrated in the context of difficult problem-solving tasks associated with crisis management or battlespace understanding. This technology will demonstrate the revolutionary impact of extensive quantities of formally represented knowledge applied to automatically interpret incomplete and uncertain situation data.

Genoa: The purpose of the program is to develop collective reasoning tools. The Genoa process collaborates and shares information from the analysis and policy maker communities. Genoa has four technologies: knowledge discovery, structured augmentation, corporate memory, and virtual collaborative environment.

Active Templates: Active Templates focuses on problem-solving methods using a spreadsheet-like interface to build simple reactive planning systems to handle real-world activation. The goal of the Active Templates program is to automate military operations, maintaining a causal model of the situation and providing incremental payoff as new automated functions are added and the causal model is improved.

Command Post of the Future: The goals of this program are to increase the speed and quality of command decisions and enable more effective dissemination of commands and smaller, more mobile command structures. The decision cycle includes situation assessment, course of action development, detailed planning, and execution. The focus of the program is visualization and HCI. The program will tailor the available information to suit the commander's situation and decision process.

Joint Force Air Component Commander: The objective of the project is to attain agile and stable control of distributed military operations conducted in an uncertain and rapidly changing environment, dramatically enhancing the effectiveness and efficiency of the JFACC. Insights gained during the initial phases of the JFACC project highlighted the need for (1) theoretical techniques and tools to support management of the dynamics in a C^2 enterprise, in addition to the planning and scheduling challenges that are prevalent; (2) a flexible modeling framework to support the development, experimentation, and proof of those techniques and tools in an environment that can be used both descriptively and prescriptively; (3) creation of new system concepts and architectures to spawn the revolution in C^2 ; and (4) development of selected prototype components to validate the new concepts and architectures.

Total Information Awareness: This program is aimed at asymmetric warfare with a transnational threat. Key components are data gathering, information discovery (model-driven search agents may be developed by industry), models and behavior (intent models, evidence models,

model-driven search agents, and inference agents), and collective reasoning (argument templates from Genoa) moving from machine to human decision making.

Intelligent Integration of Information Technology: This program is intended to provide easy access to information in the forms needed by end users and by applications. It transforms heterogeneous data collections into virtual knowledge bases by extracting, integrating, and abstracting information from the data morass. There are explicit connections with systems-oriented programs, including Battlefield Awareness and Data Dissemination, the Advanced Logistics Program, and Advanced Technology Demonstration.

Control of Agent-Based Systems: This is a project to develop the technology for controlling and coordinating large collections of autonomous software agents. The focus is long term.

High Performance Knowledge Bases: This program produces the technology needed to allow system developers to rapidly (within months) construct large (100K to 1M axiom/rule/frame) knowledge bases for military applications. Research focuses on the three procedural steps: building foundation knowledge, acquiring domain knowledge, and structuring for efficient solutions.

Image Understanding: This is a long-term program to develop computational theories and tools for integrating sensor data at multiple wavelengths. It includes automated and semi-automated exploitation tools, and the Semi-Automatic Imagery Intelligence Processing System for rapid processing of synthetic-aperture radar data sets. A security architecture is an integral part. Field experiments are included.

Planning and Decision Aids: This is a long-term research project that includes aspects of data integration, but with the primary focus on helping commanders make plans and reach decisions in real time, in rapidly changing and ambiguous environments. It is combined with an AFRL project under the acronym ARPI.

7.1.5.2 Defense Research Laboratories

7.1.5.2.1 Army Research Laboratory

Communication and Network Systems Division: This provides basic and applied research to enhance existing system concepts or to develop new system concepts involving battlefield telecommunications and information distribution.

Computer Systems Technology Division: This provides basic and applied research to provide the Army with the technology base for increased automation of land warfare material, information processing and presentation, intelligent systems, and simulation.

Information Processing and Presentation: This provides focused research to support evolving Training and Doctrine Command doctrine and addresses battlespace visualization, soldier-centered computer interface, software engineering, and intelligent agents.

Telecommunications and Information Distribution: This program focuses on wireless digital communications in an integrated architecture, joint warfighting communications, multimedia

communications, information distribution technology, and defense information technologies warfare for wireless nets.

Synthetic Environments: This provides focused research to create physically correct models to support the generation of realistic scenes of targets, terrain environments, cultural features, and battle damage. It also provides focused research to impact standards and protocols of the international modeling and simulation community, also physics-based models, 3-D visualization, and environmental and atmospheric modeling.

7.1.5.2.2 Electronic Systems Command

Data Mining in Text and Images: Data mining is emerging as a successful COTS analysis technology. Intelligence community analysts need to perform similar analysis on large text and image bases. The objective of this project is to develop technology to leverage existing MITRE strengths and COTS tools to support a new kind of analysis.

Video Markup and Interactive Multimedia Applications: The objectives of this program are to investigate competing interactive multimedia standards and their impact on DoD programs and to (1) develop a series of prototypes, using selected standards, aimed at interactivity with data for tactical operations, and (2) support DoD's needs for scene visualization through applications that support deployable environments.

Semistructured Data Research: The objective of this program is to assess the usefulness of emerging technology for (1) managing semistructured data (that is, data that has some structure but is difficult to describe with an explicit schema), and (2) integrating semistructured and structured data sources.

Goal-Driven Control of Autonomous Agents: This program is to develop an approach to building computationally inexpensive software models that exhibit robust decision-making behavior in combat simulations. This will facilitate the simulation of large, autonomous synthetic forces.

Simulation Modeling of C² Information Warfare Defense and Attack: The goals of this project are to provide empirical data on the dynamics and evolution of the information security and attack relationship and to demonstrate the feasibility of modeling evolution in complex multi-agent systems.

Adaptable Real-Time Distributed Object Management: The purpose of this project is to investigate and demonstrate advanced adaptation techniques for customizing shared (common) software infrastructure components to meet unique program requirements for building distributed, real-time, fault-tolerant C² systems without redesigning or recoding.

Migrating Legacy Systems Using Component Frameworks: The purpose of this project is to understand how to migrate legacy systems to use component frameworks (JavaBeans with CORBA and ODBMS).

Collaborative Infrastructure for the C² Enterprise: The objective of this project is to enable an environment that facilitates enterprise-wide collaboration by addressing technical issues,

capturing collaborative requirements, conducting experimentation, capturing lessons learned, and transitioning technology to the GOTS and COTS communities.

Collaborative Virtual Workspace: The objective of this program is to build a collaborative environment to enable geographically and temporally dispersed collaborators to co-locate virtually for sharing of information and applications.

Collaborative Mission Applications: The goal of this project is to examine innovative concepts for multisource intelligence analysis and product generation in a collaborative environment.

Tactical Mobile Mesh Networks: This project is developing techniques that solve the three biggest problems in tactical wireless networks: topological instability, link unreliability, and changing bandwidth.

COTS Wireless Communications: The objective of this program is to demonstrate the ability of (1) cellular mobiles and bases augmented with custom frequency translator appliques to operate in military frequency bands and (2) a digital cellular mobile augmented with a custom adaptive interference mitigation appliqué to operate in military interference environments.

Nomadic Networking and Quality of Service in Emerging Networks: This research project proposes and analyzes mechanisms for assuring quality service under dynamic conditions imposed by nomadic networks and users. Research areas in this project are maintaining multiyear, multilocation programs in nomadic networking and assuring quality of service in emerging networks.

Dynamic Configuration Management for Wireless Mobile Environments: This project investigates, develops, and demonstrates techniques that can manage and process the information required to automatically generate, dynamically update, and distribute, IP router configuration files to platforms and apply the Theory of Constraints within wireless mobile environments.

Information Fusion Environments: This project's goal is to develop an environment for the development, testing, and invocation of user-defined multisource or multimedia information fusion strategies. This would be achieved by providing systems for interaction with and visualization of these strategies and by applying fusion modeling techniques to several Air Force problems.

Advanced C² Investigation: This project aims to develop a software architecture for C² decision support and battlefield visualization by prototyping C² decision-support applications within a framework to refine or validate software architecture and by leveraging the existing C² techbase.

Battlespace Visualization: The objectives of this program are to (1) explore advanced HCI issues for the C² platform, and (2) investigate collaborative visualization and interaction with battlespace information.

Human-Computer Interaction: The objectives for this program are to (1) explore the relationship between human-human and human-machine interactions in an instrumented multiparty environment, (2) capture interactions among humans, data resources, intelligent multimodal

participants, (3) log interaction for data collection, training, evaluation, and (4) exploit data collection and machine learning to enable more productive human-machine interactions.

Language Processing for Intelligence Applications: The overall objective of this project is multilingual language processing technology with wide scope of application, including robustness despite free variability of open-source texts, and self-extensibility to break the knowledge acquisition bottleneck.

Tailoring Information Extraction Systems: The goal of this project is to port IE systems to new applications rapidly (in days to hours).

C² Protect Laboratory: The objectives of this project are to develop and promulgate techniques to address Information Warfare (IW) concerns prevalent within Air Force environments. Specifically, this project is working to (1) develop advanced tools for network discovery, policy verification, and decision support, (2) assess emerging intrusion detection and vulnerability assessment technology, and (3) develop/enhance rapid response capability for field demonstration and test.

Security for Mobile Agents: The objectives of this program are to (1) understand security issues and requirements of mobile agent systems, (2) develop techniques to secure mobile agent systems, and (3) investigate the use of mobile agents in information warfare.

Efficient Public Key Certificate Verification: The objective of this project is to identify efficient methods for certificate verification in web environments by focusing on the enhancement of the certificate revocation process.

Virtual Information Services: The objectives of this project are to create a secure framework to share business critical information via the global internet and to create a suite of new information services for DoD.

Multimedia Computing: The objective of this program is to establish knowledge of multimedia information processing and integration of associated enabling tools and technologies. This project aims to establish the growing area of multimedia information retrieval, foster connections with leading research institutions, create novel content-based video analysis algorithms, disseminate expert knowledge, and influence intelligence and C² activities.

Knowledge Management Systems: The purpose of this program is to transition processes and technology that improve the organizational learning process, including the following: (1) identification and creation—executable knowledge is created or identified; (2) diffusion—effective knowledge transfer through dialog or externalization or internalization; (3) integration or modification—integrating or adapting new knowledge; and (4) action—convert knowledge into behavior.

3-D Information Uncertainty Resolution and Cooperative Information Management: The objectives of this program are the coding and presentation of information, particularly imperfect (“uncertain”) information, to aid the user in processing and consuming information, even if this presentation is an abstract representation of the battlespace.

“Smart” Yellow Pages for Open Source: The objectives of this project are to (1) develop and provide tools to facilitate knowledge discovery and generation for sharing across the enterprise or communities of interest, (2) focus on information organization and integration technologies that link information across disparate environments, and (3) emphasize personal and workgroup information exploitation tools.

Scalable Text Summarization: The objective of this program is to address important technology gaps in text summarization. With the growing emphasis on full-text searching, information overflow problems are worsening. Text summarization can help by reducing the reading time in information retrieval tasks. It can also help in characterization of information content in data mining applications.

7.1.6 System Prototyping for Defense Applications

7.1.6.1 DARPA

7.1.6.1.2 Information Systems Office

Joint Logistics Advanced Concept Technology Demonstration (ACTD): The purpose of the Joint Logistics ACTD is to develop and integrate web-based logistics joint decision support tools into the GCSS. The Joint Logistics ACTD will make it possible for the warfighter to determine and evaluate logistics support in terms of unit mission capability instead of the traditional measures of tons of equipment and number of people moved. It will also give the warfighter or logistician the ability to quickly develop and evaluate alternative logistics concepts to support the warfighter’s possible courses of action. The Joint Logistics ACTD will also demonstrate the means to monitor the execution of logistic operations in a visualization-rich environment that supports a fused picture of the battlespace.

Joint Task Force—Advanced Technology Demonstration: The program seeks to provide “anywhere, anytime” information support to a deployed joint task force commander in battlefield operations and to support applications for use in disaster relief and law enforcement operations. It is developing a supportable, portable C⁴I software technology base for JTF crisis management, planning, and execution.

Battlefield Awareness and Data Dissemination: This demonstration includes information management aspects that appear relevant to JBI input, development of push/pull technology for manipulation and management of multiple data sets, and expansion of bandwidth (by factors of 10 to 100) accessible to tactical end users. It has explicit support from the Intelligent Integration of Information Technology program.

7.1.6.2 Defense Research Laboratories

7.1.6.2.1 Air Force Research Laboratory

Collaborative Enterprise Environment: The aim of this program is to reduce time and cost involved in developing, testing, and fielding new weapon systems through comprehensive virtual simulation. It is a distributed virtual environment that supports the collaborative use of

engineering, cost, and multiple levels of user interaction models to support design and utility testing. It emphasizes co-evolution of system artifacts and work processes.

Broadsword and C² Link: This systems development program exploits COTS web-based technology to enable rapid query and retrieval of information from diverse and distributed databases in support of intelligence operations. The system will identify sources of information and will broker access on behalf of the user to multiple classified databases defining a process approved by the National Security Agency.

Configurable Aerospace Command and Control: Emphasis is split among information management tasks, the development of collaboration and visualization tools, and training. Conducted in collaboration with AC²ISRC, C²BL, and C²TIC, this is a 5-year spiral development program with planned demonstration activities on an 18-month cycle. It will be interoperable with TBMCS and synchronized with Defense Information Infrastructure Common Operating Environment (DII COE) evolution. A base technology program in DARPA ITO (Intelligent Collaboration and Visualization) may be related but with a very long-term focus.

Joint C⁴ISR Architecture Analysis/Planning System: This system will facilitate comparing, contrasting, and integrating C⁴ISR architectures. It will allow users to easily access, display, share, and manipulate C⁴ISR architecture information through a DII COE distributed and networked system.

Command Post of the Future: The aim of Command Post of the Future is to double the speed of command post decisions using half the staff. 3-D visualization techniques are featured, together with natural language processing technology. One rubric is to increase the provision of information by exception.

7.1.6.2.2 Office of Naval Research

Quick Set: This system provides multimodal user interfaces and a multi-agent software environment to allow rapid, collaborative. It has been used to rapidly develop a diverse range of applications, ranging from a simulation setup and control system, a force laydown development system, and a database retrieval system.

7.1.6.2.3 Electronic Systems Center

Global Broadcast Service: The goal of this program is to create an architecture to support efficient dissemination management and effective attention management and to construct a framework for organizing new technologies into solutions that bring user attention to important information based on time, location, situation, schedule, priorities, needs, and available communications systems.

7.2 Gap Analysis

This section compares the technical architecture developed in the body of this report with the current and near-time expectations for the development of commercial technology. The shortfall between commercial technology and desired JBI functionality suggests the strategy for further R&D investment by DoD.

7.2.1 The Gap Between Commercial Technology and the JBI Technical Architecture

Virtually all of the core technology needed for the JBI can be purchased from the commercial sector today. The challenge is in developing the Defense-specific systems on top of this technology. An example is fusion support, which is specific to Defense needs and systems and is unlikely to be commercially supported. Another is the need for automatic target recognition—the ability to rapidly and reliably recognize specific targets within complex, high-resolution radar and electro-optical imagery.

It should be possible to buy all JBI hardware (computers, input/output devices, storage), networking (routers, fiber optics, satellite systems), and display technology components (projection displays, head-up displays) from the commercial sector.

It should be possible to construct the JBI on top of commercial technology for enterprise application integration (that is, connectors, publish and subscribe platform, etc.). This middleware is being developed to integrate numerous legacy business applications, allow business to data-mine its often-duplicative databases, and provide real-time control of business processes. These developments are being fielded in ever increasing numbers, and into increasingly complex business situations, requiring thousands of transactions per second and integrating hundreds of legacy databases in near-real time.

It should be possible to construct the JBI on top of commercial technology for data management, archiving, and data warehousing. Recent deployments of products like Oracle 8 have provided the capability to archive and interrelate data in as many ways as can be imagined. What is required is the ability to provide multilevel security on top of the commercial capabilities. While some limited capabilities have been demonstrated, it is essential that DoD continue research efforts and influence commercial products. The JBI will certainly require data inputs that span the range of classification, and will need to serve users at all levels of security classification.

It should be possible to construct the JBI on top of commercial technology for speech understanding. There are numerous corporate initiatives to improve the HCI, and speech understanding is one of the areas that has received considerable attention and investment. Bill Gates has stated that this will be a key area for Microsoft in the future for interfacing to Windows.

It should be possible to construct the JBI on top of commercial technology for visualization and virtual reality. The interactive computer gaming industry is estimated to be in excess of \$7 billion annually. This has exceeded the movie entertainment annual market. Together, however, the investments in virtual environments, visualization engines, and individual visualization environments will provide most, if not all, of the visualization capabilities required by the JBI.

It should be possible to construct the JBI on top of commercial technology for collaboration and group interaction. The power of group collaboration software is expected to increase dramatically to support highly dispersed businesses and business-to-business interactions on a global scale. As bandwidth availability increases, the ability to create virtual presence will

further improve the collaboration possible through advanced collaboration software. This effort will be fueled as well by efforts to improve the “Internet chat room” experience.

Commercial technology is focusing on methods and toolkits for more rapid deployment of integrated applications. This is useful for Defense systems as well.

Commercial technology is focusing on methods to handle scaled-up processing demands. Moore’s Law appears to be viable for a least another decade, after which feature size will eventually begin to place restrictions on unlimited processing growth. It is expected that these limits too, however, will find themselves broken by entrepreneurs who understand that enhanced processing will enable further capabilities in commercial software and interactive gaming and visualization yet undreamed of. This is useful for Defense systems as well.

Longer-term (beyond 5 years) commercial technology is interested in developing and using intelligent agent technology. It does not know how to do this today. This is an important area for continued basic and applied research investment by Defense.

In general, the JBI as it involves, depends on a higher level of intelligent, goal-oriented processing within the JBI to attain intelligent behavior on behalf of users. This will require further R&D funded by Defense, with ultimate transition into the commercial world. A key aspect that must be pursued by DoD is network-aware agent behavior. Much of the JBI will occur in bandwidth-limited or -transient bandwidth environments, which may not be the norm for the less-mobile-, commercial market, where even mobile users have significant fixed supporting infrastructure. Commercial technology will not meet traditional Defense security needs. Commercial focus is on authentication and customer privacy. Commercial cryptographic technology is sufficient. Support for multilevel security or coalition security models in commercial systems is unlikely. This is an important area for continued basic and applied research investment by Defense. Information assurance and information survival programs are very important for the JBI.

Commercial technology *may* not meet Defense reliability and robustness needs. Commercial focus is on speed of deployment, not utility (“never fail”) functionality of the result. This issue is critical in real-time weapon systems, where positive control is essential. Building reliable “good enough” systems from inherently unreliable underlying systems and information sources is an important area for continued basic and applied research investment by Defense.

7.3 Recommendations for Investment Strategy

- Buy, don’t build. Said another way, the key challenge is not to create the JBI system from scratch, but to develop it by integrating existing commercial products and services, tailored to the specific DoD mission application. This is very analogous to the commercial business market, wherein middleware vendors mold their software tools and those of other vendors to provide an integrated solution to meet the clients’ needs. This process invariably includes integration with existing legacy databases and changes to business processes—both of which must occur, for strictly adding new technology to old processes will not create the revolution envisioned for the JBI.
- Increase DoD investments in the technology base of network-aware intelligent agents, information assurance and survival ability, robust system design (error-tolerant, not necessarily error-free code),

and HCI beyond speech. These efforts are essential to implementation of the JBI and are not expected to be forthcoming from the commercial market. One key aspect of the intelligent agent market that will greatly benefit the JBI is the ability of agents to communicate with each other. By being both bandwidth-aware and able to communicate, they can, through their behavior, significantly affect the availability of data to disadvantaged users and the resultant network loading. Unless DoD takes a lead in this area, these capabilities will not be assured, and the negative impact to the JBI could be significant.

- Invest in DoD-specific fusion applications to ride on the commercial middleware. Specifically, increase efforts in multisource fusion technologies, aiming to get close to the source, rather than attempting to fuse distilled reports that have lost significant context in their distillation. The advent of increased processing power and increased bandwidth makes this a possibility. As sensor systems expand and the flood of data increases, this becomes an extremely important issue if sensor overload is to be avoided. Human analysts simply cannot keep up with the real-time sensor exploitation requirements of the JBI. While the Dynamic Database program at DARPA is working this issue, funding levels for the program are far short of what will be required to achieve the full vision of automatically fused information from all battlefield sensors.
- Continue to invest in core technologies—for example, terabit networking—well beyond the next generation.
- Increase investment in technologies that support defense C² applications—in particular, adaptive, intelligence processing, security appliques, robust system architecture, and wide-area distributed processing. The need to create a joint schema wherein each of the Service C² applications can communicate and share data will be an essential first step to the JBI. Use of XML as a key enabler for this interface should be vigorously pursued. Establishing a standard and consistent schema (also referred to as an ontology or logic) for defining or describing various data types—for example, Tank, M1A1; Vs Tracked Vehicle, Armored, U.S.; M1A1 Tank. This “wrapper software” will be critical to the level of integration and data sharing expected.
- Invest in software-programmable radios. DoD must be able to interoperate at all levels—between Services, between platforms, and with allies or coalition partners and commercial systems. While it is possible that the commercial sector will invest in this technology, the drive for competitive advantage and the implications for processing power and cost may preclude this technology in the near term. In addition, requirements for DoD-specific waveforms, jam resistance, environment-adaptive waveforms (for example, assured connectivity even at reduced bandwidth), and waveforms with low probability of intercept necessitate DoD investment irregardless. To facilitate the integration with legacy radio networks (which are imbedded in all DoD weapon systems), it is essential that software-programmable radio become the standard. This will allow for integrating mixed forces in contingency deployments and long-term upgrade to weapon systems—a capability not supported (or necessarily desired) in the commercial world. Software-programmable radios essentially become the “legacy software wrapper” equivalent to provide interfaces to and between legacy weapon systems and those of the future.
- Disinvest in research in component- or object-oriented systems, application integration, and databases.

7.4 Summary

This chapter provides roadmaps for many key JBI technologies. A view of commercially available middleware, consisting of four backplanes (web infrastructure, systems management and monitoring, security, and foundation infrastructure) is presented. The capabilities of commercial middleware products are projected, with significant gains expected by 2004. Developments in networking technology are also examined in some detail here. Finally, an analysis of the gap between commercial technologies and JBI needs leads to the investment strategy recommendations found in the previous section. More details on technologies needed to build the JBI, especially those to be managed by DoD, are provided in the next chapter.

Chapter 8: Building the JBI

8.0 Introduction

This chapter provides a strategy that brings JBI technology to the Air Force and DoD. It discusses near-term concept validations, far-term research agendas, benefits offered to business processes, and organizational issues surfaced by JBI technology. The organization of the chapter is as follows: Section 8.1 lists this study's core recommendations. Section 8.2 recommends near-term YJBI concept validations. Section 8.3 recommends longer-term research programs. Section 8.4 discusses business processes, and Section 8.5 presents specific executable recommendations for building the JBI.

Building the information-centric JBI will not be simple. Folding it into the mainstream of Air Force C² will be even less so. However, near-term prototypes need *not* be built from scratch. Inexpensive, COTS-based JBI experiments should be used. These experiments will permit warfighters to evaluate JBI technology and designers to determine the appropriate technical base for the JBI. While commercial components make up the underpinnings of the JBI, the military has unique needs, including information assurance and fused information products. Gaps between commercially available components and JBI platform architecture must be filled with DoD research initiatives. This chapter outlines a strategy for these initiatives.

The JBI is driven by e-commerce web technology. E-commerce is expanding at breakneck speed, possibly exceeding \$1 trillion by 2002.²⁶ Thus, evolutionary development is critical to maintaining information superiority with the JBI or any other web-based system. Appendix D provides the study team's recommendations for the JBI evolutionary spiral. That appendix also describes metrics appropriate to evaluating JBI spiral increments.

8.1 Study Core Recommendations

The core recommendations resulting from this study are

- Immediately fund concept validation prototypes. Each partially complete prototype (YJBI) will employ an object-based common representation. Each will include an object repository and object-based force templates. Low-cost experimental prototypes will demonstrate a service, such as publishing information from a YJBI client²⁷ to the object repository. The YJBI experiments will support development of business rules to go with new information systems. YJBI experiments should make heavy use of commercial products. These may include products inappropriate for later operational designs for reasons of security or other factors.
- Establish an initial information management cadre to work operational business processes as part of the YJBI development.

²⁶ "Intel CEO Says the Web Wave Is Just Beginning," <http://www.techweb.com>, 6 October 1999.

²⁷ A JBI client is any software application that makes use of JBI platform services to publish, subscribe, or otherwise interact with JBI information objects.

- Fund research programs within Service laboratories and DARPA to support advanced JBI platform concepts. These include issues such as common representation, military information assurance, and information fusion.
- Initiate spiral development of an operational JBI. Initiation should occur after designers and warfighters agree on functionality and after YJBI concept validations have established credibility.

Since JBI technology is inevitably part of many future DoD systems, Air Force operators must think through the following issues:

- Impacts on C² business processes
- Implications in store for warfighting staff (within C² centers)

A new reality is that information has become as important as physics to the Air Force. Even in today's mini-wars, information-centric issues such as threat location and trusted target identification dominate military business practices. They have risen in importance at precisely the same time as the World Wide Web, with its promise of click-through access to anything.

The Web is clearly in the thoughts of warfighters. For many reasons, including ones just noted, the Web will ultimately revolutionize C² business practices. The JBI of this report is an early conceptual step in the migration of Air Force C² applications to information-centric web operation. Issues remain, such as the security of distributed and mobile code, the cost of migrating existing C² systems, and the need to fit military webs to evolving business processes. Since these issues will ultimately be solved, they should not be allowed to impede immediate progress.

The JBI is one of many possible vectors for Air Force C², albeit one broadly consistent with onrushing trends in commercial technology. In DoD, developmental vectors for the DII COE, GCCS, GCSS, and TBMCS are realigning toward the Web. Similar web-based systems will be cheaply available to U.S. adversaries from global vendors.

For all these reasons, if the U.S. goal is information superiority, there is no option but to plunge into the issues of web-oriented C². The issues are more than technical. The Air Force and DoD must think carefully about the benefits JBI technologies bring to their business processes. Additionally, they must consider revising C² staff organizations if they expect to optimally employ information systems in warfighting situations.

The spiral acquisition model must be used for the JBI. There are many reasons for this (see Appendix D). There are commercial product solutions to satisfy portions of the JBI vision. As far as possible, "buy, don't build" acquisition is preferable. Spiral acquisition will enable developers to synchronize prefabricated COTS components with other pieces of the JBI. The spiral model permits developers to link

- COTS components—for example, middleware and web-based search engines
- GOTS C² systems and architectures
- Common representation standards for C² information exchange

Given its many external interfaces, careful analysis is critical to information assurance, quality of service, and other desired attributes. The study team recommends initiating system analyses of the long-term architecture of the JBI platform. It does *not* recommend that such architecture investigations impede rapid YJBI prototyping efforts, since early prototypes are crucial to the warfighter and provide architectural validations.

Because early prototypes are crucial to the warfighter, the study team recommends that business rules be developed in conjunction with the YJBI system. Furthermore, the business rules employed by users of the JBI prototypes must push the envelope of information usage. That is, while the business rules must evolve with the prototype system, they must be built to encourage users to think not only of the support the YJBI is giving now, but the possible capabilities that the next JBI could give. Those developing the business rules must be current operators who are encouraged to ask, “What if?” and, “Can the JBI do this?” These questions must be asked at several levels. At a low level, an individual wants to know how better to perform one set of tasks, while at a higher level system-wide business rules must be questioned and improved. Once a set of business processes (that is, the methods of interaction of people and the JBI to accomplish tasks) is developed, it can be modeled and improved as described later in this chapter.

The study team recommends commercial product evaluation supporting the JBI. As the commercial market expands for enterprise integration and web technology, so will opportunities for improving the JBI. The study team promotes a validation facility to examine new products for inclusion in the JBI spiral.

The study team recommends that the Air Force support DoD efforts to resolve longstanding issues of common data standards for C² software applications.²⁸ The prospect for common representation of JBI contents, agreed to by all Services and coalition partners, is tantalizing. Short of this, the Services must at least make standards open to each other. Organizing the diverse contents of the JBI into a common representation is a daunting and thankless task.

In conjunction with the data standardization effort described above, the study team recommends that the Air Force, and ultimately DoD, develop definitions for the attributes stored in force templates. The collection of force templates must be general enough to define information needs for a wide variety of units. At the same time, the number of force templates must be minimized to ensure maintainability.

Finally, the study team recommends strong Air Force participation in information assurance research that permits JBI technology to operate while under attack. Distributed components—for example, web servers, CORBA, and agents—are critical to future C² systems. DARPA and others have large research efforts in information assurance. The study team recommends supporting these efforts with Air Force research funds. The team *does not* recommend allowing the youthful state of information assurance technology to inhibit progress toward a JBI. Effective defensive technology will eventually result from research, likely in forms that permit the full

²⁸ Openness and pairwise integration is used today. The promise of the JBI will be met only if progress is made in setting data standards that become widely used.

range of JBI platform services. For now (as with *all* DOD information systems), systems must make do with less security than would be preferred.

8.2 YJBI Concept Validations

Funds should be made available immediately for rapid, low-cost YJBI prototypes. Early prototyping efforts will use a hybrid of commercial software, existing GOTS components, and custom components. Rapid prototypes form the basis for immediate concept evaluations. Taken together, these efforts are an initial step in spiral development of fielded JBI products (JBI-1, JBI-2, ...).

The spiral development of the JBI must go hand in hand with development of the information staff, both professionally and technically, and with the development of the JBI business processes. The C² business processes are tightly linked to the information systems in use. These business processes must evolve in spirals with the JBI. In addition, new business processes should be evaluated using process models, data models, workflow models, and simulation models.

At least one YJBI could be fielded as a Category 2 experiment at JEFX 00. *Major long-term architecture concerns such as information assurance should not be allowed to slow down early JBI prototypes, limit their highly experimental nature, or expand costs.* The goal should be to quickly get to demonstration, experimentation, and concept validation.

Rapid prototypes need only exhibit partial coverage of JBI core platform services: a skeletal common representation, at least one C² client application exhibiting publish and/or subscribe, a force template to describe the interaction between a client unit and the JBI, and an interaction capability. For example:

- Integration—commercial enterprise integration technology
- Common Representation—XML documents
- Force templates—represented as XML documents
- Software rule base—limited publish and subscribe represented as web scripts and as scripts interacting with commercial enterprise integration products
- Interact capability—commercial web browser
- Client application—any potential JBI client (ISR planner, fusion engine, etc.)

Rapid prototypes should be designed to clarify the vision of USAF warfighters, joint-Service users, and C² system design specialists. The focus of these first JBI prototypes should be firmly on *inexpensive evaluation and idea-generation*. Architecture studies conducted in parallel should focus on downstream functionality for spiral development.

8.3 Long-Term Research and Development

Several important long-term research areas are listed below. They *are not* listed in priority order; all are critical to the success of the JBI. They describe new (and refer to current) joint research efforts involving AFRL, DARPA, and other organizations:

- Agents and the integration of agents (DARPA lead)
- Information assurance (DARPA lead)
- Information fusion, common representation, and scripting (AFRL and DARPA)
- Acquisition and storage (DARPA /Digital Library/AFRL lead)
- Effective interfaces (AFRL lead) that
 - Adapt to the user, device, bandwidth
 - Are easy to use by the complete range of JBI users
 - Can be integrated over multiple display (visual, auditory, tactile) and control (voice, manual, gesture) modalities

In addition, there are areas where commercial agendas will be the primary drivers, moderating the need for DoD assistance at the research level. These include

- JBI platform services (commercial lead)
- Immersion environments and visualization (commercial lead)

8.3.1 Agents

Agents (mobile code entities) offer a new computing model for C² system development. Agents have been used for applications as diverse as information search, filtering, message routing, and network control. Commercial development focuses largely on e-commerce applications such as web search. The DARPA CoABS program has established much broader objectives. CoABS envisions a library of agent components, agent markup languages, and interoperability infrastructure to enhance agent (and legacy) communication, providing an information web serving the JBI as well as other systems. The study team recommends that the Air Force participate. In addition, transition of DARPA agent technology to AFRL has begun, and the study team recommends that high priority be given funds for this transition.

8.3.2 Information Assurance

This study's work on the JBI gave little attention to security issues. In part, this was because such issues are extremely broad and require a dedicated SAB study to pursue. However, security is a pressing need of the JBI, as for any element of the C² enterprise. Information attacks are of particular concern to the mobile code present in many web-based systems, including the JBI. Unfortunately, the momentum of the World Wide Web has carried all systems well beyond the ability of current defensive technology. As with any new technology, defense lags attack. The JBI, *like all DoD enterprise systems*, will continue to evolve while the young field of information assurance strives to catch up.

The study team recommends that the Air Force give strong support to information assurance research programs, such as those at DARPA. The study team especially recommends research relevant to distributed component architectures, such as CORBA, Enterprise JavaBeans, and agents. The following comments focus on three areas:

- Intrusion detection
- Automated response selection
- Multilevel security

Intrusion detection. In complex systems like the JBI, security vulnerabilities will always exist. While current signature-based²⁹ systems can identify stereotypical “cracker” fingerprints, recent research³⁰ indicates nonsignature methods may perform better against novel attacks. Current methods for “locking all the doors and windows” with signature methods should be augmented with new nonsignature methods—for example, ones that address the security of mobile code such as Java.

Response selection. Once intrusion is detected, rapid response is required to maintain system availability. Almost all attacks, whether they are manual or automated, occur at *execution speed*. Therefore, real-time attack assessment and response selection requires an approach that also runs at execution speed. This implies that research must develop control technology for automated response selection. DARPA is pursuing research in these areas, and the study team strongly endorses joint Air Force funding and AFRL participation. New theoretical models of information systems under attack³¹ are also required for these areas to progress adequately, and the study team recommends that AFRL support such work.

Multilevel security. Readers will be aware that DoD has not made great progress in multilevel security after many years of effort. Regardless, the study team recommends that AFRL participate in such efforts, given the flexibility that multilevel security would afford systems such as the JBI.

8.3.3 Fusion, Common Representation, and Scripting

Information fusion and representation methods specific to military operations are not areas that benefit from commercial research. The study team recommends them as critical areas for continuing AFRL and DARPA investment, since it will require breakthrough technology to make significant advances.

²⁹ Intrusion detection technology falls into two broad categories: signature methods and nonsignature methods. Signature methods involve such approaches as filters placed in the information stream, while nonsignature techniques include such approaches as complex sensors implanted deep within enterprise software.

³⁰ Robert Durst, Terrence Champion, Brian Witten, Eric Miller, and Luigi Spagnuolo, “Testing and Evaluating Computer Intrusion Detection Systems,” *Communications of the ACM*, 42:7, pp. 53-61.

³¹ Nong Ye, Joseph Giordano, and John Feldman, “CACS—A Process Control Approach to Cyber Attack Detection,” to appear in *Communications of the ACM*.

The essential areas of interest are:

- Universal and readily accessible information taxonomies for JBI data (for example, the SCR of the JBI repository described in an earlier chapter)
- New information integration techniques that extend the classic definition of fusion (for example, plan-to-track integration for execution management)
- JBI scripting language to support information fusion (and other JBI components such as agents and templates)

JBI common representation. The study team recommends that AFRL carry out research on JBI information representation. Part of this work concerns the adaptive integration of disparate ontology from multiple publishers. The JBI repository will take a form similar to the SCR described in Chapter 5. Since autonomous clients bound only loosely to the JBI feed the SCR, research is needed to determine the best means of universally codifying client products for subsequent manipulation by JBI consumers.

New information fusion techniques. The JBI requires a revision to the classic four-level fusion model (see Chapter 5 of this report). The study team recommends that the model be extended to include an informational as well as a physical battlespace. The traditionally separate research communities for fusion, planning, and information assurance must work together to develop the basis of a new form of common operating picture, broader model of fusion, and new tools for integration.

As the information battlespace intrudes more on the turf of traditional warfighters, the notion of the common operating picture will be further extended to include information warfare. Traditional notions of collection management are similarly extended; they now include software-intrusion detection probes as well as ISR platforms, and informational entities as well as the physical entities.

JBI-related research should focus on automating Level 2 and Level 4 processes, particularly those emphasizing the atypical nature of JBI fusion and collection management. At Level 2, the processing steps include inferences as well as mathematical calculations, and military knowledge is often used. Level 4 deals with collection management, exploitation management, security management, and other high-level (meta) processes. Among other things, it closes the loop with external JBI publishers, satisfying information needs for improving the common representation.

The study team recommends that the Air Force transition relevant parts of the DARPA Dynamic Database and Adaptive Image Manager programs, which address classic issues associated with fusion and ISR collection management and are important to the JBI. The study team recommends participation in the DARPA Information Assurance program to gain access to models of the information battlespace, as well as to understand the software sensors needed for situation awareness. Finally, the study team recommends new AFRL programs that combine and extend these ideas.

JBI scripting language. Small JBI scripts (fuselets) are the mechanism by which the JBI integrates information in its common representation. Sophisticated fusion effects can be achieved

very simply by a fusion component library attached to the JBI scripting language. Similar component libraries to support agents and active templates will markedly improve the power of the JBI system for everyday users. The study team recommends that AFRL carry out the work needed to make such a language available, drawing on work available in the commercial and research community. This work should be mindful of the security problems associated with some potential choices (for example, JavaScript).

8.3.4 Acquisition and Storage

Automatic data capture has two aspects: (1) determining what information is needed and (2) automatically capturing it. Templates address the first of these, and new information-centric business practices the second. *Templates* are complete (but empty) data schemas describing information needs. Attaching software to basic templates produces *active templates*, able to guide users through the process of filling empty data schema. The study team supports DARPA's Active Templates program for this technology.

New wearable computers and wireless networks will allow the JBI to capture many information products useful to the JBI *in situ*. The challenge is not technical, but managerial: U.S. forces need new *information-centric business practices* that support use of this technology. The study team recommends that the Air Force develop and apply these business practices.

Information extraction is the ability to convert raw unstructured information to the JBI SCR. *Shallow extraction* refers to simple information such as physical entities, numbers, or events. For example, raw numerical logistics data may be captured *in situ* as described above and inserted directly into precast templates. *Deep extraction* is needed for information such as complex military scenarios. For example, a commander's mental situation model could be extracted via semiotic screen representations and converted to JBI objects. The study team recommends that AFRL develop an information extraction program supporting the JBI.

Web-based retrieval is important to the JBI, for which is envisioned online access to all information associated with the C² enterprise. An essential question will be how to avoid overwhelming busy C² staff with irrelevant information pulled up during retrieval. The JBI offers to bring the C² information portfolio online, but must avoid a high volume of irrelevant digital retrievals. The unstructured nature of much C² information makes the problem harder: how can one retrieve a particular target from the many Predator videos in the JBI object repository?

Many web-oriented search engines function by measuring text-based similarity between query and content. Others, such as <http://www.google.com>, are based on measures of mass-consumer interest: how many consumers have pointed their browsers at a particular website? Neither approach may be powerful enough to dip into the JBI SCR and retrieve a threat network lying close to the path of an in-bound F-117.

Issues of retrieval are being addressed in the Digital Library initiative. The study team recommends AFRL participation in these projects.

Advanced data storage. The JBI will require persistent data stores. New database technology is needed, including advanced query methods and distributed or heterogeneous processing. This

will offer the JBI platform new data modeling and routing methods and better information integration tools. DARPA funded early research in these areas and successfully transitioned to commercial markets. However, transition within DoD has languished. Except for the Army Research Laboratory, few DoD transitions have taken place. The study team encourages AFRL involvement to ensure that such technology is available to Air Force C² systems.

8.3.5 Effective Interfaces

The goals of the JBI are to provide the right information to the right person at the right time in the right media, right language, and right level of detail. Several technologies support these goals.

Dynamic user modeling enables the JBI to tailor information based on personal preferences, task, and device. These models filter and format information for users as well as manage workflow. The DARPA Command Post of the Future program is experimentally evaluating this technology. The study team recommends that AFRL monitor these experiments. If they prove successful, AFRL should extend the technology for use in Air Force C² centers.

Collaborative information integration is enabled by three technologies: (1) dialog management enhancing communication among groups separated by distance and time, (2) context understanding to track location and roles of users, and (3) shared collaborative applications. The study team recommends continued funding for AFRL and DARPA programs developing these technologies.

8.3.6 Commercial Research and Development

There are some important areas where commercial initiatives predominate. For these, the Air Force should focus on integration rather than research. An example is *JB I platform services*, provided by the commercial middleware technology described in Chapter 6. Another is *immersive environments and visualization*, described in Chapter 4, where entertainment markets drive research.

8.4 Business Processes and the JBI

The study team recommends that the Air Force consider the impact of JBI technology on its business processes. These considerations are critically important to the long-run success of modern Air Force warfighting concepts, including reachback, strategy-to-task, EBO, sensor-to-shooter, and precision attack. Like any modern organization, the business practices of the Air Force C² enterprise are tightly linked to its information systems. Driven by e-commerce technology, Air Force enterprise systems like the JBI will evolve at a fast pace. Business processes must evolve with it.

The JBI can provide solutions to many new operational issues. Where will warfighters find the ability to dynamically replan the ATO, link execution management to planning, and reach decisions faster than their opponents? This report is predicated on the assumption that the JBI will enable fresh and perhaps unexpected operational concepts through new means of information

retrieval, management, sharing, interpreting, and discovery.³² The study team believes that these new capabilities will advance sufficiently in the near future to merit substantial attention by Air Force C² operators.




Figure 21. Operational view of dynamic retargeting




Figure 22. Information-centric view of dynamic retargeting

Such comparative analyses should be supplemented by hands-on evaluation. In the end, only actual use will decide whether these technologies can provide the returns envisioned in this report. Experiments conducted with YJBI prototypes will have great bearing on how soon JBI

³² See the Stanford Digital Library Project web pages, including <http://www-diglib.stanford.edu/diglib/pub/>.

technologies are accepted by mainstream C^2 users. Thus, emphasis is placed in this report on spiral development and its tight coupling of software development with iterative user evaluation of both the software and the associated business processes.

Success in future wars will require new abilities to rapidly access large amounts of shared information. It follows that business processes must evolve if they are to optimally utilize these new information technologies. Thus, it is this study's recommendation that AC²ISRC and others begin evaluating how business processes might be altered to reflect the capability of JBI-like systems. For example, Figures 21 and 22 compare the operational picture of dynamic replanning with the equivalent JBI object flows. The figures shows how the JBI replaces the traditional, linear approach to dynamic retargeting with an approach in which the entire planning status is constantly changing depending on inputs from a variety of sources. The user sees the current situation, knowing that it may need to change based on information received in near-real time. Clearly, this new approach changes the workflow for the planner.

Finding the best business processes for humans and machines may require new tools for modeling Air Force C^2 business logic. To date, process models and data models have received primary emphasis. *Process models* such as integrated definition represent activity flow in C^2 processes. However, these often only codify the flow of control among existing software applications and current C^2 business rules. *Data models* are justifiably receiving great attention across DoD.³³ The study team supports Air Force participation in DoD-wide efforts toward data standardization.

Perhaps closer to operator issues lies the less-explored territory of *workflow models* and *simulation models*. The former are used to represent the human actions that are part of C^2 , and the latter enable dynamic "what-if" analyses involving entities and actions within the C^2 enterprise. U.S. forces are far from having complete workflow and simulation models for battle management. In fact, U.S. forces compensate for lack of such models by carrying out experiments such as JEFX and by using spiral methods to modulate software development. The study team recommends that the Air Force invest the time and effort required to develop such models and to use them to facilitate the evolution of business processes.

8.5 The JBI Roadmap

The JBI must be developed as a major weapon system. Figure 23, a roadmap for development of the JBI, is modeled after an aircraft development roadmap. Like aircraft development, building the JBI requires significant long-term planning and commitment of resources. The roadmap presented here lacks funding information, which must be programmed over the coming year, but the roadmap should support the development of accurate cost estimates. The JBI roadmap shows a variety of fronts on which development must start immediately: testbed standup, business rules,

³³ For examples see the DISA DoD Data Administration page <http://www-datadmn.itsi.disa.mil/mission.html>, and the DII COE Shared Data Engineering page <http://diides.ncr.disa.mil/shade>. However, the JBI will be more useful if it rests on a deeper information representation (described in the Input section) than is the case with these laudable DoD efforts.

COTS evaluation, common information representation, fuselets, advanced data retrieval, and agent technology.




Figure 23. JBI Development Roadmap

During YJBI development, an initial capability should be achieved after the testbed has been operating for 2 years and lessons have been learned from two JEFXs. This allows for development of the initial business rules and most of the common information representation, as well as integration of appropriate COTS technologies. The testbed will stay in operation as YJBI business rules are refined.

The final capability of JBI-1 depends on an early start on research in fuselets, data retrieval, advanced fusion, and agent technology. In Block 10, military requirements from all Services must be well understood before a JBI-1 design is pursued. The JBI-1 technical architecture development commences after the military requirements analysis, two rounds of JEFXs, and the associated lessons learned from the YJBI. Force templates, the key to integrating units into the JBI, will be developed in Block 10, as will fuselets. Information assurance is a critical concern for the JBI, and this technology will be built into JBI-1 during Block 10. The roadmap shows an initial capability for Block 10 in late 2003, with refinements made for at least a year after that.

JB1-1 Blocks 20 and 30 consist of the integration of technologies that are less mature today, namely advanced fusion, agent technology, and user modeling. These technologies aid decision making, automate higher-level reasoning, and optimize user interaction with the JBI. Advanced fusion should mature by mid-2005 to give Block 20 an initial capability, while Block 30 will achieve an initial capability in 2007.

JB1-2 development starts in 2007. By then, many of the aforementioned technologies will be mature and available in COTS form. JB1-2 development will consist of integrating these well-understood technologies and pushing progress of appropriate leading-edge technologies.

8.6 Specific Executable Recommendations

This chapter describes in some detail the approach that should be taken to start building the JBI. This section highlights the key recommendations stemming from this SAB study. The technical recommendations described in this chapter include:

- AFRL and AC²ISRC should stand up a JBI testbed
- AFRL and AC²ISRC should start development of low-cost JBI prototypes immediately
- ESC should evaluate relevant COTS systems for inclusion in the JBI and develop the JBI technical architecture
- AC²ISRC should develop the military requirements for C² information integration
- DISA and ESC should jointly develop the common data representation and information for force templates
- DARPA and AFRL must initiate or accelerate long-term research in the following areas:
 - The advanced JBI platform and technical architecture
 - Advanced fusion concepts
 - Information assurance
 - Agent-based technology
 - Advanced data survivable systems
 - Active networks
 - Dynamic user modeling

(This Page Intentionally Blank)

Chapter 9: Overall Recommendations

9.0 Introduction

Most of the recommendations in the previous chapter were of a technical nature. However, building the JBI is much more than a technical challenge—it is a leadership challenge as well. Senior leaders must support development of the JBI, and they must ensure that doctrine, organization, training, materiel, and people move forward with the technology. To that end, this chapter describes management-related steps to be taken to ensure that the JBI goes from the vision described herein to reality.

9.1 Commit to the JBI as a Major Weapon System

A system of the JBI's size and scope cannot come to fruition through small experiments in a handful of different labs. The current Service efforts at improving C² systems must be marshaled together under one JBI joint program office. With central management, redundancies can be eliminated, and efforts can be steered to development of critical JBI technologies. Long-range planning based on the roadmap presented in Chapter 8 will ensure that all aspects of the program are addressed. To ensure that the JBI is adequately funded, the Services must earn the support of civilian leaders, especially Congress.

9.2 Create an Information Staff Function

As shown in Chapter 3, the information staff is a critical cog in JBI operations. The Services must recognize the importance of recruiting, training, and retaining warriors who can manipulate information and information systems. The information staff is not composed of communications officers. While their training includes topics such as web interfaces, use of objects, modeling, scripting, computer security, information operations, and networking, members of the information staff must have expertise in functional areas like operations, logistics, and intelligence. Retention of personnel with high-tech skills is a challenge for all Services. A viable, rewarding career must await those who choose to be infowarriors, or retention problems will continue. Without the staff to manage it, the JBI will be a hollow shell.

9.3 Develop New Concepts of Operations at AC²ISRC

As the JBI develops, the business processes that accompany it must also be developed. Operators must work hand in hand with JBI developers to ensure that the business processes embedded in the JBI are viable. At the same time, these operators must not constrain themselves to doing business as usual. They must find new ways to take advantage of the services provided by the JBI—constantly asking, “What if?” and, “Can we?” As business processes change, a significant investment in training will be required to get all users up to speed.

9.4 Define Common Information Representations Led by the Assistant Secretary of Defense for Command, Control, Communications, and Intelligence (ASD-C³I)

The JBI's SCR and ontology must be built on a common information representation. Attempts to enforce common representation between Services have failed to date. This is a huge task that, when completed, will reap huge rewards. It must be pursued vigorously now. A common representation will enhance interoperability and ensure that the JBI builds a consistent picture of the battlefield situation. The sooner a common representation is agreed on, the fewer new applications will require wrappers to transform data into the JBI's representation.

9.5 Reinforce DARPA R&D Investment for JBI Technologies

Following the recommendations in the previous chapter, DARPA will be the lead agency for much of the noncommercial research needed to support the JBI. DARPA must give this research high priority, redirecting current efforts as needed. Obviously, this research must be backed with appropriate funding.

9.6 Focus AFRL, Other Service Research Labs, and Battlelabs on Evaluating and Applying Commercial Technologies for the JBI

While DARPA manages longer-term research, Service labs should immediately start evaluating commercial technologies for integration into the YJBI. XML and other commercial web technologies make an ideal foundation for the YJBI. Commercial enterprise integration technology can be used to make the different pieces of the JBI work together. Interaction technologies are rapidly maturing due to the interest in commercial markets. The JBI must leverage as much commercial technology as possible.

9.7 Create the JBI Testbed Now for JEFX 00 Participation

JBI development must start now so that a meaningful prototype can be used at JEFX 00. JEFX 00 provides operators with their first chance to evaluate the earliest incarnation of the YJBI in an operational environment. If this opportunity is missed, valuable operator input based on the JEFX environment will be delayed for at least a year. Furthermore, the deadlines associated with exercises give developers an incentive to quickly mature the JBI into useable form.

9.8 Link the JBI Testbed to Other Service Efforts in Digitized Battlefield and Network-Centric Warfare

The JBI must be joint. The Air Force by itself cannot develop the JBI into a system with the capabilities described in this study. The other Services must be involved as early as possible in JBI development, preferably under the direction of a joint program office. A significant step toward a fully joint JBI is to link existing and near-future Army and Navy systems with the JBI. A common data representation won't be found on all these systems in the near term. However, linking these systems can be a significant test for commercial enterprise integration technology in the role of connector or wrapper between the different Services' systems.

9.9 Evangelize the CINCs

As the primary customers of the JBI, the CINCs must support its development and deployment. Like operators at all levels, they must be given a chance to ask, “Can the JBI do this?” and say, “What I really want is this.” When they understand, and ultimately experience, the information superiority the JBI provides, they too will champion the cause of Building the JBI.

9.10 Summary

Building the JBI is much more than a technological challenge. It goes hand in hand with a sea change in the way information is used and, ultimately, in the way wars are fought. While the technologists build the system, senior DoD leadership must lay the groundwork for its deployment. This chapter has described the management-related steps that must be taken to ensure that all facets of the JBI, from doctrine to organization to logistics to materiel to people, evolve with the technology.

(This Page Intentionally Blank)

Acronyms and Abbreviations

ABCS	Army Battlefield Control System
ABIS	Advanced Battlespace Information System
ABL	airborne laser
AC ² ISRC	Aerospace Command and Control Intelligence, Surveillance, and Reconnaissance Center
AC ² TIC	Aerospace Command and Control Training and Innovation Center
AcT	Active Templates
ACTD	Advanced Concept Technology Demonstration
AF	Air Force
AFMC	Air Force Materiel Command
AFOSR	Air Force Office of Scientific Research
AFRL	Air Force Research Laboratory
AI	artificial intelligence
AIM	Adaptive Image Manager
ALP	Advanced Logistics Program
AOR	area of responsibility
API	applications programming interface
APORTS	aerial ports
ASD-C ³ I	Assistant Secretary of Defense for Command, Control, Communications, and Intelligence
ASR	automated speech recognition
ASTOR	Airborne Standoff Radar
ATD	Advanced Technology Demonstration
ATM	Asynchronous Transfer Mode
ATO	air tasking order
AWE	Advanced Warfighting Experiment
BADD	Battlefield Awareness and Data Dissemination
BDA	battle damage assessment
BI	Battlespace Infosphere
C ²	command and control
C ⁴ I	command, control, communications, computers, and intelligence
C ⁴ ISR	command, control, communications, computers, intelligence, surveillance, and reconnaissance
CACC	Configurable Aerospace Command and Control

CAPS	C ⁴ ISR Architecture Analysis/Planning System
CDE	Common Data Environment
CEC	Cooperative Engagement Capability
CEE	Collaborative Enterprise Environment
CINC	commander in chief
CIO	Chief Information Officer
CJTF	Commander, joint task force
CM	cruise missile
CoABS	Control of Agent-Based Systems
COE	common operational environment
COM	Common Object Model
CONOPS	concept of operations
CONUS	continental United States
CORBA	Common Object Request Broker Architecture
COTS	commercial off-the-shelf
CPoF	Command Post of the Future
CSS	Combat Service Support
DARPA	Defense Advanced Research Projects Agency
DCI	Director of Central Intelligence
DDR&E	Director, Defense Research and Engineering
DB	database
DIA	Defense Intelligence Agency
DII	Defense Information Infrastructure
DISA	Defense Information Systems Agency
DMIF	Dynamic Multi-user Information Fusion
DoD	Department of Defense
DOTML-P	doctrine, organization, training, materiel, leadership, and people
DTD	Document Type Definition
EFX	Experimental Force Exercise
EJB	Enterprise JavaBeans
EO	electro-optical
FAA	Federal Aviation Administration
Gb	Gigabits
GBS	Global Broadcast Service
GCCS	Global Command and Control System

GCSS	Global Combat Support System
GOTS	Government off-the-shelf
HCC	High Confidence Computing
HCI	human-computer interaction
HLF	high-level fusion
HPKB	High-Performance Knowledge Bases
HTML	Hypertext Markup Language
HUMINT	Human Intelligence
IC&V	Intelligent Collaboration and Visualization
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IMINT	Imagery Intelligence
IO	Information Operations
IP	Internet Protocol
IPNG	Internet Protocol, Next Generation
IPv6	Internet Protocol, Version 6 (Next Generation)
IPA	imagery product archive
IS	information superiority
ISO	Information Systems Office
ISP	Internet Service Provider
ISR	intelligence, surveillance, and reconnaissance
IW	Information Warfare
J2EE	Java 2 Enterprise Edition
J2ME	Java 2 Platform, Micro Edition
JBIP	Joint Battlespace Infosphere
JBIP	Joint Battlespace Infosphere platform
JCAPS	Joint B24
JCS	Joint Chiefs of Staff
JCTN	Joint Control and Tracking Network
JEFX	Joint Expeditionary Force Exercise
JFACC	Joint Forces Air Component Commander
JFLCC	Joint Forces Land Component Commander
JFMCC	Joint Forces Maritime Component Commander
JFC	joint force commander
JMS	Java Messaging Service

JSTARS	Joint Surveillance, Target, and Attack Radar System
JTF	joint task force
JTIDS	Joint Tactical Information Distribution System
JTN	Joint Targeting Network
JTS	Joint Technical Architecture
JV 2010	Joint Vision 2010
Kb	Kilobits
KQML	knowledge query manipulation language
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
MASINT	Measurement and Signals Intelligence
Mb	Megabits
MRC	Major Regional Conflict
MTI	moving-target indicator
NASA	National Aeronautics and Space Administration
NCA	National Command Authority
NCW	network-centric warfare
NGI	Next Generation Internet
NIMA	National Imagery and Mapping Agency
NIST	National Institute of Science and Technology
NRO	National Reconnaissance Office
NSA	National Security Agency
NSF	National Science Foundation
ODBMS	Object-Oriented Database Management System
OODA	Observe, Orient, Decide, Act
OS	Operating System
OSD	Office of the Secretary of Defense
P&DA	Planning and Decision Aids
QOS	quality of service
RAP	Repository Access Protocol
R&D	research and development
RDBMS	Relational Database Management System
RDF	Resource Description Framework
RFC	Request for Comment
RKF	Rapid Knowledge Formation

ROE	rules of engagement
RSVP	Resource Reservation Protocol
SAB	Scientific Advisory Board
SAM	surface-to-air missile
SAGE	System for Automatic Graphic Expression
SCR	structured common representation
SGML	Standard Generalized Markup Language
SIGINT	Signals Intelligence
SMTP	Simple Mail Transport Protocol
SOF	Special Operations Forces
SQL	Structured Query Language
STAR	System Technology for Advanced Resolution
TBM	theater ballistic missile
TBMCS	Theater Battle Management Core Systems
TCP/IP	Transmission Control Protocol/Internet Protocol
TEL	transporter-erector-launcher
TIA	Total Information Awareness
TMD	Theater Missile Defense
TOS	type of service
TPED	tasking, processing, exploitation and dissemination
TRADOC	Training and Doctrine Command
TRANSCOM	Transportation Command
TTP	tactics, techniques, and procedures
UAV	unpiloted aerospace vehicle
URL	uniform resource locator
USAF	United States Air Force
VPN	Virtual Private Network
VRML	Virtual Reality Markup Language
WMD	weapons of mass destruction
XML	Extensible Markup Language

(This Page Intentionally Blank)

Appendix A: Terms of Reference

BACKGROUND: The USAF Scientific Advisory Board (SAB) 1998 Information Management to Support the Warrior Study recommended an approach to combat information management called the Joint Battlespace Infosphere. The Joint Battlespace Infosphere will support the warfighters at all echelons in obtaining the needed information to conduct combat operations. The Joint Battlespace Infosphere uses a publish and subscribe concept that allows the warfighters instantaneous access to information needed without being overwhelmed with data. It also permits the automatic transformation of data to create higher levels of knowledge. The recommendations are consistent with previous Defense Science Board and Air Force SAB studies.

Global information is the keystone to enabling Global Engagement and *Joint Vision 2010*. The operational requirement is to develop a flexible system that supports long-range strike missions, and major theater wars, as well as regional small-scale conflicts involving all the other Services and coalition forces. Commanders need to be able to tailor the information management system to their own specific needs while still having access to the global information. Additional study is necessary to further define the Joint Battlespace Infosphere and its implementation.

STUDY PRODUCTS: Briefing to SAF/OS and AF/CC in October 1999. Publish report December 1999.

CHARTER: Continue to study combat information management and develop specific recommendations that support the implementation of a Joint Battlespace Infosphere.

1. Continue to assess the commercial research in information management technology so that the advances may quickly be applied to combat information management systems through a spiral development process.
2. Identify approaches for creating combat information management systems and for developing rule-based information distribution processes.
3. Identify interoperability issues to determine how to support the Service-unique and coalition (to the extent possible) combat information requirements when operating in a joint and/or combined environment.
4. Investigate and document where DoD resources need to be applied to support the military unique requirements in combat information management.
5. Develop an implementation plan including key demonstrations, to define a spiral development process aimed at achieving an operational capability (perhaps incrementally) as soon as technologies are available.

(This Page Intentionally Blank)

Appendix B: Study Members and Affiliations

Study Chair: General James McCarthy, USAF, Ret
Olin Professor of National Security, U.S. Air Force Academy, CO

General Officer Participant: Maj Gen Jerry Perryman, Commander
Aerospace Command and Control Intelligence, Surveillance, and Reconnaissance Center

Study Members

Affiliation

Input Panel

Dr. Charles Morefield (Panel Chair)
Prof. Edward Feigenbaum
Prof. James Hendler
Dr. Robert Miller
Maj Gen Richard O'Lear, USAF, Ret

Alphatech
Stanford University
University of Maryland/DARPA
MITRE
Oracle/Lockheed Martin

Manipulate Panel

Dr. Robert Sproull (Panel Chair)
Dr. Barry Leiner
Dr. Scott Renner
Dr. William Rouse
Mr. Thomas Saunders
Dr. Northrup Fowler

Sun Microsystems
Corporation for National Research Initiatives
MITRE
Enterprise Support Systems
MITRE
AFRL/IFT

Interact Panel

Dr. Valerie Gawron
Dr. Duane Adams
Dr. Llewellyn Dougherty
Mr. Scott Fouse
Dr. Bob Eggleston

Calspan, an operation of Viridian
Carnegie Mellon University
Raytheon Systems Company
ISX
AFRL/HECA

Commercial Panel

Prof. Randy Katz
Mr. Troy Crites
Mr. Carl Kessler
Mr. Sean Rice
Dr. Steve Wolff

University of California, Berkeley
Sparta, Inc.
IBM
The Boeing Company
Cisco Systems

Joint Aspects Panel

Vice Adm David Frost, USN, Ret (Panel Chair)
Lt Gen Carl Franklin, USAF, Ret
Mr. Ed Mahen

Frost and Associates, Inc.
Falcon Associates
SRI International

Government Advisors

Brig Gen Ben Robinson, USAF
Maj Steve Jenkins, USAF
Dr. Kevin Kreitman
Mr. Brian Sharkey
Mr. James Shaw

AF/XOC
National Security Space Architect
The Aerospace Corporation
DARPA
32 AOS/AOW

Support

Study Executive Officer:
Study Technical Writer:
Panel Executive Officers:

Capt D. Brent Morris
Maj Jason Moore
Maj Laura Olsen
Capt David Gaines
Capt Matt Yocum

AF/SB
USAF/DFCS
AC²ISRC/C2PF
USAF/DFM
USAF/DFEM

Appendix C: Study Meetings

Kickoff Meeting

22–23 February 1999

Hurlburt Field, FL: C² Training and Information Center

Information Gathering Meeting

4–5 March 1999

Charleston, SC: SPAWAR System Center-Charleston, U.S. Army Digital Integration Lab (DIL),
Command and Control Unified Battlespace Environment (CUBE)

Information Gathering Meeting

9 March 1999

Stanford, CA: Stanford University

Information Gathering Meeting

12 March 1999

Washington, DC: ANSER

Information Gathering Meeting

17 March 1999

Pittsburgh, PA: Carnegie Mellon University

SAB Spring Board/Study Meeting

24–27 March 1999

Colorado Springs, CO: Space Command

Information Gathering Meeting

12–13 April 1999

Boston, MA: Sun Microsystems Laboratories, Inc.; The MITRE Corporation

Information Gathering Meetings

15–16 April 1999

Mountain View, CA: GTE

Moffett Field, CA: NASA-Ames Research Center

Palo Alto, CA: Tibco

San Jose, CA: BEA Systems

Sunnyvale, CA: Vitria

Redwood Shores, CA: Oracle

Redmond, WA: Microsoft

Information Gathering Meeting

23 April 1999

Nellis AFB, NV: Green Flag Exercise

Information Gathering Meetings

17–20 May 1999

Chapel Hill, NC: University of North Carolina

North Reading, MA: Lotus

Bohemia, NY: Logicon-Northrop Grumman

Murray Hill, NJ: Lucent Technologies–Bell Laboratories

Princeton, NJ: Sarnoff Corporation

Washington, DC: DARPA and the National Science Foundation

Springfield, VA: Autometric Inc.

College Park, MD: University of Maryland and Army Research Laboratory

Information Gathering Meeting

24–25 May 1999

Washington, DC: Lockheed Martin and Alphatech

SAB Summer Session

14–25 June 1999

Irvine, CA

Appendix D: Spiral Development

The study team recommends that the JBI be developed using a spiral process. This process encompasses user evaluation and technical architectures defining the underlying substrate. It also encompasses processes for updating objectives, negotiating increment scope, evolving components, managing risk, and measuring outcomes. A spiral approach is required because the JBI is a hybrid of three technologies: COTS (middleware), GOTS (C² applications), and custom software (SCR). Furthermore, JBI COTS components are evolving extremely quickly. A spiral will permit the JBI to maintain pace with its evolving commercial components.

Table 4. *Example of a JBI spiral*

Spiral Phase	Objective	Task	System Interfaces	Technology
1	Define common object types	Create foundation of the digital library. Define nature, relationship, dissemination of information objects specific to the military mission	Internal to JBI platform	Object modeling, schema development
1.5	JBI client system integration (for example, ISR or fusion component)	Define integration methods for client under test, publish thread, no subscribe	E.g.: RAOC databases	XML, message-oriented middleware, web server, other interfaces
2	Integration of publish/subscribe	Modify client to publish and subscribe	Single system integration	Enterprise integration technology
2.5	Legacy system integration, first increment	Define bidirectional integration methods for components under test	System of systems	Enterprise integration technology
3	Improve system performance	Metric examination, feedback incorporation, implementation exercises	System of systems	Business practice evaluation

D.1 Spiral Increments

The spiral model endorses the concept of *build a little, test a lot*. The JEFX is an excellent venue for testing. Others include the Battlelabs, AFRL, ESC CUBE, and Langley, Hurlburt, and Nellis AFB C² facilities. JBI tests may include participation by all of the locations noted.

Development spirals are designed to answer many questions. For example: What constitutes an acceptable COTS foundation? What capabilities should receive priority? Table 4 offers an example event sequence that might be applied to the JBI spiral.

Current policy supports tailoring. However, most acquisition strategies employ a waterfall model, not a spiral. To accommodate differences between the two acquisition models, the study team recommends a *spiral-to-increment* approach. The goal should be to deliver small,

compatible increments providing additional capability every 6 to 18 months. Each increment should address a specific mission segment and deliver measurable benefits not dependent on future increments.

Challenges for hybrid systems such as the JBI include

- Modification of COTS components to interoperate with custom and GOTS software
- Maintenance, licensing, and ownership issues
- Synchronization independently evolving COTS/GOTS/custom components
- Interoperability across spiral increments.

Setting the stage for later increments should be as important as deploying the current increment. The JBI segments should be as brief as feasible. Small increments reduce risk, minimize schedule delays, avoid cost overruns, and keep the focus on “crawl-walk-run” evolution.

D.2 Spiral Precepts

There are several precepts that guide spiral acquisition, design, and deployment of the JBI. The spiral uses a *standards-based* approach, since it focuses vendors on interoperability. It also protects technology investments, since commercial research is attracted to widely used standards. Avoiding proprietary components makes it easier to deploy middleware. Potential standards for the JBI are numerous: Java and Enterprise JavaBeans, CORBA, the various IP standards (TCP/IP, HTTP, XML, MIME, etc.), and others. Many are now (or likely to become) part of the DII COE.

A *server-centric* approach exploits the web model. Essential processing occurs on servers, while web browsers provide access. Local processing is done via Java applets, which can be controlled, updated, and downloaded by a server. This significantly lowers deployment and maintenance costs. It helps avoid large client software footprints. Unstable clients have less impact, and emerging device types (for example, wearable computers) are more easily supported.

To *leverage core systems*, designers must take full advantage of legacy, not rewriting software or forcing its premature obsolescence. This may be accomplished by means such as wrapping legacy programming interfaces, publishing legacy data as XML, or adding metadirectory³⁴ front ends to legacy databases.

Scalable systems are built to run on heterogeneous software and hardware platforms, permitting easy addition of processing power. It is difficult to accurately predict compute workloads for web-based systems, so the ability to quickly adapt to new loading patterns is essential. Enterprise JavaBeans are an example of a scalable technology.

Evolutionary development is *quick to deploy*—it starts with simple cases, enables early concept validation, and adds functions only at stable increments. It focuses on software that is *easy to use* during installation, configuration, administration and operation. A web browser client is an example of an easy-to-use approach since it follows well-known user metaphors.

³⁴ The term *metadirectory* was coined by the Burton Group to describe tools that integrate legacy directories.

Manageable software achieves system continuity and availability through system management tools (for example, SNMP, Tivoli). This allows system administrators to monitor run-time performance and get immediate notification of component failures.

D.3 Operational Test and Evaluation of the JBI System



Figure 24. JBI metrics

Figure 24 provides an overview of metrics that may be used at spiral increments. Quantitative metrics will play an important role during concept evaluation, providing a means to distinguish among alternative JBI designs. During concept validation, focus should be on the unique nature of the JBI, which requires unique metrics. Ideally, metrics will distinguish between two JBI designs offering different implementations of equivalent functions. Several metric concepts are described below.

System availability. The JBI is a transaction-oriented system that will be realized by networked servers. Metric thresholds evaluating the stability of complex systems of interacting servers with very large transaction loading must be set appropriately: failures occur and overall reliability is far from perfect. The Schwab online electronic trading company provides an example.³⁵ Schwab's site has as many as 75 million page hits per day. With as many as 60,000 simultaneous users, Schwab claims 99 percent availability.

The information quality ratio. Mechanisms for delivering information are often oblivious to end users' needs (for example, the mechanisms may involve irrelevant message traffic or high

³⁵ From "As E-Commerce Surges, So Do Technical Problems," Matt Richtel, the *NY Times on the Web*, 21 June 1999.

density of irrelevant icons on screens). A significant benefit of the JBI is better control of extraneous traffic and better integration of objects ultimately presented to users. Information quality ratios measure the number of pertinent information objects presented to the total in the object repository. This will be important in evaluating JBI query, publication, and subscription services, since poor designs lead to a confusing flood of data.

Information latency. Future C^2 systems will have substantially more communications bandwidth than today. This will not necessarily translate to increased presentation speed. Information latency benchmarks the ability of the JBI to accelerate the observe, orient, decide, and act loop. Measurements should indicate when objects are presented to JBI subscribers (rather than when the originating event occurred).

Collaboration coverage. This metric measures the degree of increase in collaborative workflow. Fully coordinated decisions are a result of relevant objects moving among decision makers connected to the JBI. The degree to which the JBI supports task coordination, within a meaningful C^2 timeline, is a measure of coordination effectiveness. The collaboration coverage metric measures the quantity of JBI object flow among collaborating users within the time they use to reach a decision.

Business process improvement. The JBI will improve current business processes. Examples include continuous ATO generation or dynamic ISR collection management. Business process metrics capture the ability of the JBI to support meaningfully new C^2 functions. For example, there is a significant benefit for the JBI to acquire point of use data, such as when weapons are loaded onto aircraft. The immediate availability of such information to the planner will accelerate the observe, orient, decide, and act loop. It is also a component needed for continuous ATO production.

Report Distribution

Headquarters Air Force

SAF/OS	Secretary of the Air Force
AF/CC	Chief of Staff
AF/CV	Vice Chief of Staff
AF/CVA	Assistant Vice Chief of Staff
AF/HO	Historian
AF/ST	Chief Scientist
AF/SC	Communications and Information
AF/SG	Surgeon General
AF/SF	Security Forces
AF/TE	Test and Evaluation

Assistant Secretary for Acquisition

SAF/AQ	Assistant Secretary for Acquisition
SAF/AQ	Military Director, USAF Scientific Advisory Board
SAF/AQI	Information Dominance
SAF/AQL	Special Programs
SAF/AQP	Global Power
SAF/AQQ	Global Reach
SAF/AQR	Science, Technology and Engineering
SAF/AQS	Space and Nuclear Deterrence
SAF/AQX	Management Policy and Program Integration

Deputy Chief of Staff, Air and Space Operations

AF/XO	DCS, Air and Space Operations
AF/XOC	Command and Control
AF/XOI	Intelligence, Surveillance and Reconnaissance
AF/XOJ	Joint Matters
AF/XOO	Operations and Training
AF/XOR	Operational Requirements

Deputy Chief of Staff, Installations and Logistics

AF/IL	DCS, Installations and Logistics
AF/ILX	Plans and Integration

Deputy Chief of Staff, Plans and Programs

AF/XP	DCS, Plans and Programs
AF/XPI	Information and Systems
AF/XPM	Manpower, Organization and Quality
AF/XPP	Programs
AF/XPX	Strategic Planning
AF/XPY	Analysis

Deputy Chief of Staff, Personnel

Office of the Secretary of Defense

Other Air Force Organizations

U.S. Army

U.S. Navy

168

Initial Distribution (continued)**U.S. Marine Corps**

DC/S (A) Deputy Chief of Staff for Aviation

Joint Staff

JCS	Office of the Vice Chairman
USJFCOM	U.S. Joint Forces Command
J2	Intelligence
J3/5	Operations
J4	Logistics
J6	Command, Control, Communications & Computer Systems
J7	Joint Training
J8	Strategy, Requirements & Integration
J9	Joint Experimentation

Other

Study Participants
Aerospace Corporation
ANSER
MITRE
RAND

(This Page Intentionally Blank)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and manipulating the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE December 1999	3. REPORT TYPE AND DATES COVERED Final, January 1999 - December 1999		
4. TITLE AND SUBTITLE Building the Joint Battlespace Infosphere, Volume 1: Summary		5. FUNDING NUMBERS		
6. AUTHOR(S) J. McCarthy, D. Frost, R. Katz, R. Sproull, C. Morefield, V. Gawron				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AF/SB Pentagon Washington, DC 20330-1180		8. PERFORMING ORGANIZATION REPORT NUMBER SAB-TR-99-02		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) SAF/OS AF/CC Pentagon Washington, DC 20330-1670		10. SPONSORING/MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Public Distribution authorized, February 2000		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 Words) Building the Joint Battlespace Infosphere The Joint Battlespace Infosphere (JBI) is a combat information management system that provides individual users with the specific information required for their functional responsibilities during crisis or conflict. The JBI integrates data from a wide variety of sources, aggregates this information, and distributes the information in the appropriate form and level of detail to users at all echelons. The JBI was originally described in the 1998 USAF Scientific Advisory Board (SAB) report <i>Information Management to Support the Warrior</i> . The JBI is built on four key concepts: 1. Information exchange through "publish and subscribe," 2. Transforming data into knowledge via fuselets, 3. Distributed collaboration through shared, updateable knowledge objects, and 4. Assigned unit incorporation via force templates. The report provides a brief description of the JBI's capabilities and its operation. The main focus of the report, however, is the technical roadmap for building the JBI. Namely, it answers the questions "What technologies make up the building blocks of the JBI?"; "In which technologies should the Department of Defense invest?"; and "How should that investment be managed?" In order to ensure that the JBI goes from the vision described in the report to reality, the report recommends that the JBI be considered a major weapon system. In that light, a JBI development roadmap is presented.				
14. SUBJECT TERMS Battlespace Infosphere, C4I, command and control, decision support tools, combat support, data fusion, fuselet, battlespace awareness, information management, information access, information grid, information manipulation, collaboration, force templates			15. NUMBER OF PAGES 188	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT None	

(This Page Intentionally Blank)