Chapter 4 – APPLICATION LAYER INFORMATION EXCHANGE ISSUES

Digitization of Command and Control information seeks to facilitate increased operational tempo by reducing the need to slow down, or stop, in order to regain situational understanding. In order to realize this operational objective, information must continue to flow as the forces move throughout the battlespace. Since commanders may need to receive and disseminate situational awareness and execution information from any location in the area of operations, all of the sources and destinations of information should remain available to the commander at all times. The objective of managed information exchange is to support the commander's ability to execute command and control by providing a timely flow of accurate, relevant information.

Computers assist human beings by maintaining and manipulating large amounts of data that represent meaningful information to the human beings that use the computers. For information to be managed and manipulated by a computer, however, it must be highly structured. This structure typically takes the form of a structured data schema implemented in a database. The structure is important to permit the computer to deal with the data. However, the structure also serves to preserve a context for the data. Preservation of context (structure) for data is critical as the data are manipulated and exchanged.

The operational benefits of digital information technology employed in support of military command and control are derived mainly from the speed with which computers can access, retrieve, process and display structured data that have been entered into them. However, the operational benefits are also highly dependent upon the ability to exchange structured data of operational importance between computers with the speed necessary to support operational tempo.

4.1 STRUCTURED MESSAGING

Over voice radio, the standard means of communication for decades has been verbal transmission of structured military messages (example: "0 this is 6, Contact Grid 247653, six enemy tanks using routes in to the built up area and concentrating in that location oriented South. Continuing to observe. Time of contact 1430 hours. Over."). In the conversion from voice to digital technology, one option is to preserve the structured message paradigm in digital form. It has the advantage that it is a paradigm with which armies are already familiar. However, in a bandwidth-constrained data communication environment, use of structured military messages has two important disadvantages. First, a structured message must, by definition, be semantically complete (i.e., in a form sufficiently complete to be fully understood by the recipient without resort to external references). Second, the message structure itself creates a certain amount of overhead (header, field separators, etc.). The need for semantic completeness can create a tendency on the part of users to recall a previously sent message, modify the fields whose content have changed, and resend the entire message (including fields whose values have not changed). Sending of redundant information and the message structure itself combine to create a significant communication overhead that can limit the throughput of operationally important new data in a bandwidth-challenged radio environment.

4.2 DATA REPLICATION

An alternative to the use of structured messages is to perform information exchange via direct database-to-database exchange. Using this approach, one or more database transactions modify attributes in relational

tables in a local database using some combination of insert, update and delete operations. The local database transactions are packaged for transmission and copied using data replication middleware to recipient nodes. At the receiving end, the database transactions are unpackaged and applied directly to the databases in the recipient nodes. Such database-to-database exchanges can be more bandwidth-efficient than structured message exchanges because they can have minimal overhead and transmit only data values that have changed.

The term 'replication' refers to the systematic propagation and maintenance of copies of data between datastores within a distributed computing environment. The field of replication has two domains; file replication and data replication. Most commercial products address one of these two but not both. File replication is concerned with the copying of complete files while data replication operates at the level of database transactions.

In database systems, a transaction is a set of database operations that must succeed or fail as a single unit. A transaction can be considered a logical unit of work that transforms the database from one consistent state to another consistent state. Data replication middleware is used to copy database transactions from a source database to multiple replicate databases. In the military context of the present workshop, the databases in question are contained in computers located in stationary or mobile command headquarters, in military vehicles, or on dismounted soldiers. From the communications perspective, each of these entities is considered a node on a military command and control network.

4.2.1 Data Replication in a Bandwidth-Constrained Wireless Environment

4.2.1.1 Synchronous versus Asynchronous Replication

The distinction between a synchronous and an asynchronous communication paradigm is common to many topics dealing with data exchange. Although the topics may appear to employ separate usages of these terms, the usages have some key points in common. The fundamental characteristic of the synchronous paradigm is that all participants of the communication are engaged at the same time. Usually, synchronous communication or data exchange protocols include an immediate acknowledgement of the reception and some schema for retransmission in case of failure. If either or both of these characteristics are absent in a communication protocol, the protocol is called 'asynchronous'. 'Asynchronous' may refer to the fact that no (immediate) acknowledgement is sent back to the sender and/or to the fact that sender and receiver engage in the communication at different times.

The following paragraphs examine the use of an asynchronous communication protocol by data replication middleware, while Chapter 5 discusses use of synchronous and asynchronous protocols by middleware in general.

Data replication middleware is used to copy database transactions from a source database to multiple replicate databases. In a commercial application (e.g., a financial system), the transactions are applied synchronously (i.e., simultaneously), using a protocol known as two-phase commit. The protocol ensures that either all participating databases commit the transaction at the same instant, or none do. Use of such a protocol produces what is termed 'tight' consistency of database content. Tight consistency means that data values are consistent (i.e., the same) across source and replicate databases at all times.

The two-phase commit protocol is a two-step protocol in which each replicate database sends a positive acknowledgement back to the source database for each step. In a bandwidth-constrained environment, use of such a protocol generates such a significant communication overhead that it can seriously impede throughput of operational data. In this case, asynchronous replication must be employed. Under asynchronous replication

the transaction(s) are applied to the source database before being replicated. Replicated transactions are applied to the replicate databases sometime after they are applied to the source database. Use of asynchronous replication produces what is termed 'loose' consistency of database content[1]. Loose consistency means that modified data values always become consistent across source and replicate databases, but only after a certain time delay (latency).

In the very low bandwidth regime associated with combat net radio, where the data throughput of the radio network can be chronically less than the load offered to the network, it is necessary to relax traditional consistency expectations. In this situation, at any point in time, a percentage of data values in the source and replicate databases will be inconsistent, producing a condition that has been dubbed 'lazy' consistency. Since one can achieve neither tight nor loose consistency of full database content in this situation, the challenge shifts to preservation of consistency for those attributes judged to be of highest operational importance. This problem is inherently more challenging than the problem of achieving tight or loose consistency. In the latter two cases, the end state is inherently stable, i.e., the database content will always evolve toward a state of consistency, albeit with some delay. In the case of chronically inadequate and variable bandwidth, the end state is uncertain and unstable. Data delivery is on a 'best effort' basis and there is no assured 'audit trail' for the database changes that are propagated. A continual deterioration in the consistency of database content is likely, unless steps are taken to manage information flow in a proactive manner.

In this low bandwidth environment, the data replication and transport mechanisms must include intelligent information management protocols capable of adapting to the variable throughput of the radio network and to changing battlefield priorities.

4.2.2 Desirable Characteristics of the Data Replication Service

In his keynote address at the data replication workshop, Chamberlain [3] from U.S. Army Research Laboratory discussed characteristics that a data replication service should have to function well in the tactical radio domain. He noted that the variation in bandwidth is as important as the actual bandwidth values in determining data throughput. Zero is a valid value for throughput in this domain. A command and control node must continue to function even when the data throughput equals zero. He proposed the use of local predictive algorithms to estimate changes in data values (e.g., vehicle position) during these periods.

To be useful, the information exchange service must support rapid selection and dissemination of information and hands-off operation. According to Chamberlain, the service should have three characteristics:

1) Automatic (hands-off, context-sensitive);
2) Adaptive (context-based, responsive to changing bandwidth); and
3) Affordable (treated as part of a Distributed Computing Environment).

The goal is to balance or 'tune' information management to available network resources.

Chamberlain's proposal for what he terms 'Model-Based Battle Command' is based upon the following four principles:

1) Use data abstractions as the medium of exchange (data replication);

[1] For this to be true, it is essential that the asynchronous replication mechanism preserve transactional semantics. Database-to-database replication of transaction events as they occur in the source database generally preserves transactional semantics whereas a complete or incremental periodic refresh (snapshot) mechanism does not.

2) Control exchange by active database triggers (provides context-sensitivity since triggers link the decision to replicate to the state (i.e., value) of certain database elements);
3) Allow reasonably different perceptions of battlefield (accept 'lazy consistency' of database content); and
4) Treat database synchronization as the realistic control of differing perceptions (management of 'lazy consistency' rather than enforcement of loose/tight consistency of database content).

4.2.2.1 Network Awareness

A key element in the above approach is the ability to monitor or measure bandwidth so that synchronization efforts can be adapted to current communication resources. To achieve this objective, according to Chamberlain a communications model, incorporating information such as network performance and connectivity data, should be part of the data schema. As well, network performance statistics should be passively collected (e.g., channel access delay, round-trip time, queue delay, number of failed versus successful transmissions), analyzed, and results made available to the application through the communications model in the data schema. Chamberlain also proposes that innovative protocol mechanisms be considered to permit more direct performance feedback to the application and more efficient data transmission. Possible protocol mechanisms include: stack cognizance (sharing of information between ISO layers, e.g., transport and data link layer), 'just-in-time' packet construction (packing several small (highest-priority) application Protocol Data Units (PDU) into a Maximum Transmission Unit (MTU)), and overhearing (the direct passage between the network and transport layer to the host of packets rather than addresses).

4.2.2.2 Data Ownership

According to [4], data replication service must enforce one or more of the following three models for data ownership: single, dynamic (i.e., transferable) and shared. 'Ownership' of a given data element refers to the right to modify the value of that data element. In a single-ownership model, the originator of the data retains ownership throughout the lifetime of the data element. In a dynamic-ownership model, any given data element has a single owner at any point in time, but the ownership is transferable; the present owner of the data element may or may not be the originator of that data element. In a shared-ownership model, there is no single owner for the data element. Any participating database may initiate a change in the value of the data element. A shared-ownership model should be avoided because it maximizes the possibility of data conflicts, and resolution of data conflicts can generate an undesirable increase in network traffic. Furthermore, it is important that the system track the identity of the owner of the data element and the ownership model that applies to that data element (when more than one ownership model is employed).

4.2.2.3 Data Recovery

It is also important [4] for data retransmission and recovery to be intelligently managed. At the packet level, this means tying the number of attempted retransmissions to the assigned transmission priority for the packet and to the known current performance of the network. At the level of bulk recovery (e.g., by a node that has been disconnected from the network for an extended period), this means minimizing the size of the recovery package by being selective as to content (for example, by ignoring time-sensitive data whose operational value ages quickly). With the latter approach, by definition the content of the recovering database will never be fully consistent with that of the other databases. However, a bulk recovery may tie up the radio sub-net for minutes or tens of minutes. By limiting the content of the recovery package, one is trading off consistency of stale data in one database in favour of consistency of fresh data in all databases.

4.2.2.4 Functional Requirements – Data Replication Service

A data replication service can be considered to consist of two parts. The Replication Mechanism (RM) is responsible to determine when replication of data from a local node to other network nodes is to occur and what information is to be replicated. It also packages and unpackages data to be replicated into/from a unit called a Replication Protocol Data Unit (PDU)[2]. A PDU is information that is delivered as a unit among peer entities of a network that may contain control information, address information, or data. The Replication Transport Layer (RTL) sits beneath the Replication Mechanism and provides a transport mechanism for the purposes of passing replication data, in the form of Replication PDUs, over the communications bearer.

The combat net radio environment, when compared to wired networks, is characterized by:

- Low bandwidth (e.g., 1 Kbps – 80 Kbps);
- High and variable latency (e.g., 0.5 sec); and
- Intermittent connectivity.

For optimal performance in the tactical wireless domain, the communication protocol in the Replication Transport Layer needs to be matched to both the characteristics of the communications bearer and the information exchange requirements of the Replication Mechanism.

One of the discussion groups at the data replication workshop identified a list of key functional requirements that should be satisfied by the Replication Mechanism and Replication Transport Layer. The list was organized into three parts, namely:

1) Functions to be implemented in the Replication Mechanism;
2) Functions to be implemented through some combination of the Replication Mechanism, the Replication Transport Layer, and, perhaps, an additional application; and
3) Functions to be implemented in the Replication Transport Layer.

4.2.2.4.1 Functional Requirements

A) Replication Mechanism:

- Intelligence to determine when replication is to occur. This function must be context sensitive.
- Intelligence to determine what data is to be replicated once a decision to replicate has been taken.
- Assembly of the replication protocol data unit (PDU). If information is available on what replications a sender or senders already possess, it is possible to determine if the PDU is semantically complete.

B) Combination of Replication Mechanism and Replication Transport Layer or other application:

- Functionality must be provided to enforce and mediate dependencies of RM and RTL on other system architecture components. Examples: dependence of RM on characteristics of a particular DBMS or other application, dependence of RTL on characteristics of a particular Data Model, RM, etc.
- Data Ownership: RM and RTL must be able to identify the owner of any given piece of data at any given time. This is tied to the policy for database key management and requires decisions about the

[2] In some analyses, the Replication Mechanism is further decomposed into a Replication Agent (the component that determines when replication is to occur and what information is to be replicated) and a Replication Server (the component that packages and unpackages data to be replicated into/from Replication PDUs).

implementation level and authority structure for such management. These decisions can depend upon how the data ownership information is being used (enforcing transactional integrity, resolving data conflicts, providing traceability of ownership over time, etc.).

- Intelligence to determine the level of effort devoted to transmitting a Replication PDU across the network based on the importance of the information contained in the PDU. For example, (1) how many retransmission attempts should be made by the RTL before it stops trying to transmit the PDU and, (2) if the RTL has more than one class of transport service (e.g., 'guaranteed' and 'best effort' transport services), which one should be used for each Replication PDU? Assuming the RTL has several mechanisms to distinguish classes of service, one of the significant issues is defining the criteria for assigning a level of 'importance' to each PDU.

C) Replication Transport Layer:

- Must support an Acknowledgement Scheme (negative acknowledgement preferred; positive acknowledgement scheme is very bandwidth-intensive).
- Must supply a retransmission protocol that takes account of time-varying communications bearer performance.
- Priority Scheme: RTL must support a prioritization scheme at the PDU level that takes account of time-varying bearer performance.
- Must enforce a degree of fault tolerance and be sensitive to time-varying communications bearer performance.
- Must support fragmentation and defragmentation of PDUs, as appropriate, determined by the characteristics of the network.
- Addressing Scheme: RTL should support an addressing scheme. Choices are a broadcast communications protocol or an addressing scheme that would allow point-to-point or point-to-multipoint communications to occur between RTL peers.
- Cognition of Network Structure: RTL should be cognizant of the network structure (who is on the radio net and the status of communications transmitted from each node). This information, if available, could be used to determine what information transmitted by a remote node has not been received by the local node and vice-versa. This capability would, however, introduce overhead to the RTL communications protocol.

4.2.2.4.2 Non-Functional Requirements

- The RM and RTL must operate effectively in both the non-disadvantaged and disadvantaged communication environments.

4.2.2.5 ATCCIS Replication Mechanism (ARM)

One of the most important examples of selective data distribution in the NATO context is the ATCCIS Replication Mechanism [5]. The present section contains a brief summary of the ATCCIS replication mechanism (a fuller description can be found in Annex D). Some advantages and disadvantages of the selective data distribution model versus all-informed data distribution model for the tactical wireless domain are discussed in Section 4.2.2.6.

ATCCIS (Army Tactical Command and Control Information System) was an international programme consisting of NATO nations (although not a formal NATO effort) aimed at identifying the minimum set of specifications to be included within C2ISs to allow the automatic transfer of selected command and control (C2) data. Their objective was to develop a specification for a hardware/software/vendor-independent interoperability solution. The ATCCIS programme ran from 1982 to 2002.

The Multinational Interoperability Programme is an international programme consisting of NATO nations (also not a NATO effort) whose focus is the fielding of an interoperability solution for multinational C2ISs. In 2002, ATCCIS merged with MIP. MIP adopted the products of the ATCCIS work as the basis for direct database-to-database exchange. However, MIP also maintains a structured message exchange mechanism.

The ATCCIS concept of interoperability is based upon the automatic transfer of standardized data elements that utilize agreed and common data identifiers based upon a common data interchange model. The common data model is called the Land C2 Information Exchange Data Model.

Additionally, the ATCCIS programme developed the specification for a mechanism that will permit interoperability of automated C2ISs through the partial replication of database content. The ATCCIS Replication mechanism (ARM) is selective in:

a) The data to be exchanged;

b) The recipients of the data; and

c) The transfer facility to be used.

Under the ATCCIS concept, nations use the common data model to preserve the meaning and relationships of the information exchanged between C2ISs across national boundaries. Nations are free to develop differently structured C2 databases for national use. The ARM is used to manage the exchange of information between databases of C2ISs across national boundaries based on the common data model.

The ARM employs the concept of contracts and filters. A *replication contract* is the means for controlling (selective) replication of database changes. A contract is established between a pair of replication nodes, designated as Data Provider (DP) and Data Receiver (DR). In the contract, the DP and DR agree that the DP will provide the DR with all data that satisfies the conditions of the contract. A contract specifies a *filter* and parameter values used to set filter conditions, as well as a DP and a DR. A filter is a set of criteria applied to the instances of a database in order to reduce the total set of data selected to a subset. Examples of filter types include geographical area, time, and order of battle (organizational). The contracts enforce a 'push' model in which the only data pushed to recipient nodes are those negotiated with the recipient node under the pre-agreed contract. To modify the set of data pushed to a particular data recipient by a data provider, a filter must be applied or the contract must be modified.

4.2.2.6 'All-Informed' Data Distribution Model versus 'Selective' Data Distribution Model

Two forms of data distribution model are important for the tactical wireless domain. An 'all-informed' data distribution model is based on the assumption that there is value in maintaining synchronized data content across all participating nodes on a subnetwork, at least for an agreed subset of entities within the data model. A 'selective' data distribution model makes the opposite assumption, i.e., that only data of interest to a particular node should be sent to that node. The underlying assumption in such a data distribution strategy is that a node should only receive data from external sources that are important to its assigned role (or that are permitted by a security policy). By the very nature of this distribution model, database content will not be

consistent across nodes on a subnet. An important example of this approach, the ATCCIS Replication Mechanism, was described in the preceding section.

A selective data distribution model ensures that a user of data is not forced to deal with data that are of no interest or importance to his role, and that, in the event of an unsuccessful retransmission, no effort is expended retransmitting the data to nodes that have no identified need for it. As well, it limits the quantity of received data that must be stored at the receiver node. Finally, if a point-to-point addressing scheme is used, it can support a security policy based on selective dissemination since transmission of a data payload can be limited to specific receiver nodes. However, in the tactical wireless domain, there can be serious operational disadvantages associated with selective data exchange due to the very limited bandwidth and variable connectivity of the wireless network. The disadvantages are:

a) Does not take advantage of shared medium (radio)

Most of the time, nodes on a wireless data subnet overhear data transmissions that are addressed to other nodes on the subnet. When data replication is used as the exchange mechanism, the database changes contained in those transmissions can be applied to the database in the overhearing node, even though they are not addressed to, or directly pertinent to the role of, the overhearing node. An all-informed data exchange model exploits these overheard data transmissions to the maximum extent possible while a selective data exchange model routinely discards this 'free' information.

b) May introduce single point(s) of failure

Under selective data distribution, each node is the sole provider of data which it 'owns', and subsets of its owned dataset are shared with a set of data receiver nodes. Each receiver node is subscribed, in principle, to a unique subset of the provider's dataset (the subsets for different receivers may overlap, but are, in general, not identical). Under this scheme, a given receiver node will only have knowledge of data being shared with it, not with the other receiver nodes, since the exchange agreements are between node-pairs. If the provider node is disabled, due to enemy action or other reason, for all practical purposes each receiver node will be limited to the narrow slice of the provider's dataset received from the provider node, and the set of receiver nodes will have no simple means of recovering or reconstructing all or part of the total dataset on the provider node, if that proves necessary (certainly no mechanism efficient enough to work acceptably over a tactical radio system). Under an all-informed data distribution model, nodes on the subnet overhear and capture data transmissions, even those for which they were not explicitly identified as a data receiver. If a provider node is disabled, each node on the subnet will have a copy of the content of all (or almost all) data transmissions emanating from the provider node to that point in time recorded in its database, thereby largely avoiding the problems noted above.

c) Requires data recovery from multiple nodes

In the event that a node leaves the subnetwork for a period of time, rejoins the network, and wants to recover data that it has missed, it will have no choice but to request missing data from each data provider node with which it has an exchange contract. If there are n such provider nodes, that will require n separate data transmissions. If three nodes need to recover data at the same time, and they have exchange contracts with n_1, n_2 and n_3 provider nodes respectively, that will require $n_1 + n_2 + n_3$ distinct data transmissions. This recovery process increases the traffic level on the radio channel and wastes bandwidth. By contrast, in an all-informed data exchange model, the node(s) can recover missing information in a single data transmission from any neighbour node, since the database content is, in principle, the same across all nodes.

d) Node cannot readily assume role of neighbour node without substantial one-time data transfer from that neighbour node

If a node is required to assume the role of a neighbour node, under a selective data exchange model it would be required to download all, or a substantial portion of, the information specific to the neighbour's role from the database of the neighbour node, since it would not have information specific to the neighbour's role in its own database. If the neighbour node database is not available, due to enemy action or other reason, it would have to request recovery of those role-specific data from the *n* data provider nodes with which the original node had an exchange contract. This would require n distinct (potentially large) data transmissions. In an all-informed data exchange model, no such large data transmissions are required because, in principle, the replacing node has the same data in its database as the replaced node.

An all-informed distribution model can still enforce a security policy of selective dissemination. Tactical radio transmissions are routinely encrypted for transmission and then decrypted at the receiving end. Thus the received decrypted transmission can be assumed to have originated from an authentic source. Selective dissemination can be enforced at the software level by implementing control software in the receiving nodes to ignore a (decrypted) data transmission if a security caveat in the received transmission instructs it to do so. With this model, data will be shared on an all-informed basis across all nodes on a subnet except for instances when the security policy specifically excludes this option.

If a point-to-multipoint addressing scheme is employed, it is possible to combine an all-informed distribution model for certain types of data (e.g., reports of friendly vehicle positions) with a selective distribution model for other types of data (e.g., passage of a fragmentary order from battalion HQ to company HQs). This hybrid distribution model will permit full synchronization of database content for certain types of data, and partial synchronization of database content for other types of data. Such a hybrid model may well provide the best match between operational requirements and bandwidth utilization in the tactical wireless domain.

4.3 DATA EXCHANGE USING XML

The Extensible Markup Language (XML) is a data interchange format developed by the World Wide Web Consortium and released in 1998. Inspired by the success of HTML (Hyper Text Markup Language), the established language of the Web, it is a simplified (and therefore more useable) subset of the older, more sophisticated SGML (Standard Generalized Markup Language). As such, XML is actually a meta-language describing a family of languages, namely languages defining interoperable, text based data formats. Although the historical roots of XML are in publishing, XML is also suited to the task of unambiguously identifying complex data structures that may never be viewed or printed. The ability to represent complex data structures and the data portability derived from the interoperable nature designed into XML, have made XML the data interchange format of choice for many applications where generality and portability are important.

An XML-based document has both a logical and a physical structure. The logical structure allows a document to be divided into named units and sub-units, called 'elements'. The physical structure allows components of the document, called 'entities', to be named and stored separately, sometimes in other data files (permitting information re-use and use of non-XML data (e.g., image data) by reference). An XML processor module is used to manage entities and combine them into a single data stream for validation by a parser and for accessing by the main application.

Because XML is a meta-language, there is no pre-defined list of elements. The user may name and use elements as desired. To allow automatic syntax checking, the user may use a Document Type Definition or an

XML Schema to define the elements allowed in a particular type of document. XML enables a high degree of control over the logical document structure. Unlike HTML, XML is targeted at the definition of data structures rather than text formatting. It therefore encourages the use of names for the elements that describe the nature of the object, rather than how it should be displayed or printed. This generalized markup approach means that the information is self-describing, and so can be located, extracted and manipulated as desired.

As a general rule, if one implements a data exchange mechanism and data exchange format flexible enough to handle the general case, the price one pays is that one is required to transmit more metadata (most of which describes data or document structure) to preserve that flexibility and generality. The 'self-describing' property of XML derives from the structure metadata (tags et al) accompanying the data content. This trend runs contrary to what is required in a low bandwidth communications environment, where the objective is to maximize data content and to minimize associated metadata in a transmitted payload. One option to reduce the size of XML data is compression. A description and comparison of different XML compression mechanisms can be found in [6].

4.4 SUMMARY AND CONCLUSIONS

Operational benefits of digital information technology employed in support of military command and control are derived mainly from the speed with which computers can access, retrieve, process and display structured data entered into them. Benefits are also highly dependent upon the ability to exchange structured data of operational importance between computers with the speed necessary to support operational tempo. In an operational military environment, two forms of data exchange are employed: structured message exchange or exchange of database transactions. The latter form can be more bandwidth-efficient because it propagates data changes only. It is implemented via a data replication mechanism.

The term 'replication' refers to the systematic propagation and maintenance of copies of data between datastores within a distributed computing environment. In the tactical domain, the datastores in question are contained in computers located in stationary or mobile command headquarters, in military vehicles, or on dismounted soldiers, all operating forward of battalion and communicating via a combat net radio. The combat net radio environment, when compared to wired networks, is characterized by very low bandwidth, high and variable latency, and intermittent connectivity.

A data replication service consists of two basic components, a Replication Mechanism and a Replication Transport Layer. The Replication Mechanism is responsible to determine when replication is to occur and what information is to be replicated. The Replication Transport Layer provides a transport mechanism for the purposes of passing replication data over the communications network.

To be effective in the low-bandwidth tactical wireless environment, the Replication Mechanism and Replication Transport Layer must include intelligent information management protocols capable of adapting to the variable throughput of the radio network and to changing battlefield priorities. This adaptivity requires that the mechanisms must have dynamic awareness of the network state and battlefield state.