GRIFFIN DIRECTORY SERVICES TECHNICAL ARCHITECTURE

ACP 137

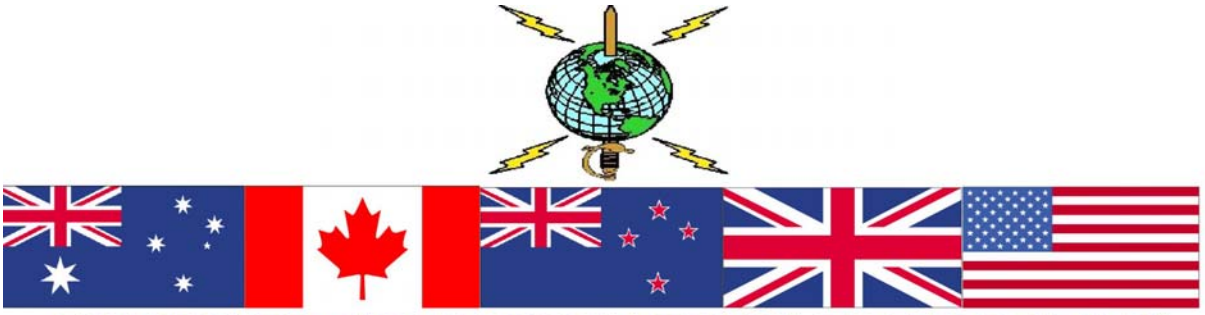

COMBINED COMMUNICATIONS-ELECTRONICS BOARD (CCEB)

OCTOBER 2007

FOREWORD

1. The Combined Communications-Electronics Board (CCEB) is comprised of the five member nations, Australia, Canada, New Zealand, United Kingdom and United States and is the sponsoring authority for all Allied Communications Publications (ACPs). ACPs are raised and issued under common agreement between the member nations.

2. ACP 137 GRIFFIN DIRECTORY SERVICES TECHNICAL ARCHITECTURE is an UNCLASSIFIED CCEB publication.

3. This publication contains Allied military information for official purposes only.

4. It is permitted to copy or make extracts from this publication.

5. This ACP is to be maintained and amended in accordance with the provisions of the current version of ACP 198.

THE COMBINED COMMUNICATION-ELECTRONICS BOARD LETTER OF PROMULGATION

FOR ACP 137

1. The purpose of this Combined Communication Electronics Board (CCEB) Letter of Promulgation is to implement ACP 137 within the Armed Forces of the CCEB Nations. ACP 137 GRIFFIN DIRECTORY SERVICES TECHNICAL ARCHITECTURE is an UNCLASSIFIED publication developed for Allied use under the direction of the CCEB Principals. It is promulgated for guidance, information, and use by the Armed Forces and other users of military communications facilities.

2. ACP 137 is effective on receipt for CCEB Nations. NATO Military Committee (NAMILCOM) will promulgate the effective status separately for NATO nations and Strategic Commands.

EFFECTIVE STATUS

Publication	Effective for	Date	Authority
ACP 137	CCEB	On Receipt	LOP/COMAG

LAST REVIEWED

Date	Authority

3. This ACP will be reviewed periodically as directed by the CCEB Permanent Secretary.

4. All proposed amendments to the publication are to be forwarded to the national coordinating authorities of the CCEB or NAMILCOM.

For the CCEB Principals

JA STOTT
Lt Cdr RN
CCEB Permanent Secretary

RECORD OF MESSAGE CORRECTIONS

Identification of Message Correction and Reference Date		Date Entered	By whom entered
Reference Date	Correction		

TABLE OF CONTENTS

TABLE OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

OVERVIEW

101. This document specifies a Griffin Directory Service (DS) technical architecture. The Griffin DS is required to meet the needs of new and evolving services on the Griffin network, including High Grade, Military Messaging (MM) via the Allied Communications Publication (ACP) 145 Gateway architecture, and its need for a Public Key Infrastructure (PKI).

102. The Griffin DS is based on the Griffin Interim DS [Ref 6], and is intended to be backwards compatible with it, until such time as the Griffin Interim DS capability is subsumed. The Griffin DS has been enhanced to allow:

a. A choice of transport mechanism.

b. Support for further attributes and objects required for support of MM.

c. Support for more complex national directory replication requirements, including large data volumes and different transfer formats.

103. In this model each nation will still share their directory information with the other CCEB nations in the form of LDIF files, although further formats may be added as national requirements change. These files will be transmitted using an allowable transport mechanism which initially includes the existing Griffin email service and the ACP 145 based MM service.

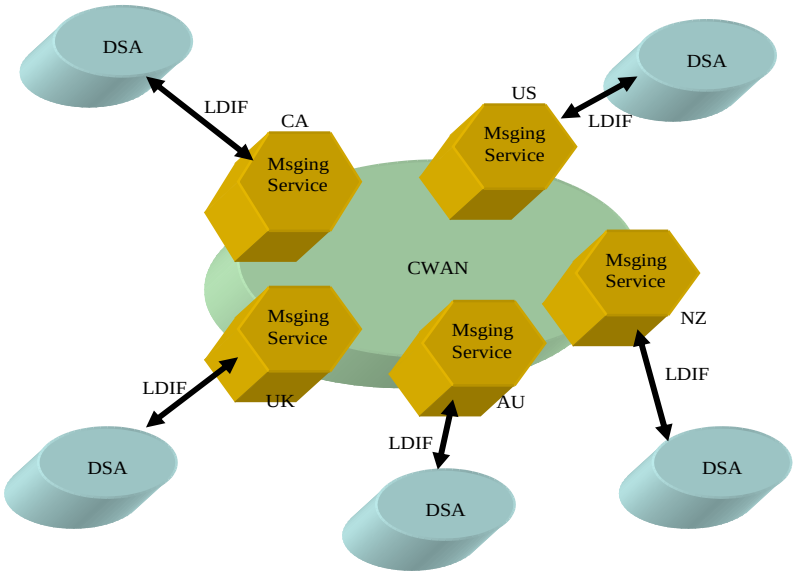

Figure 1-1 GRIFFIN Directory Service Architecture

104. Figure 1 shows the conceptual architecture for Griffin DS. It does not preclude future enhancement to support other transport mechanisms including a “Pull” capability such as download via Web pages. In order to facilitate this capability, an optional mechanism has been allowed to specify a URL from which the file may be accessed, rather than sending the file as an attachment to the message.

DOCUMENT SCOPE

105. The scope of this document is that portion of the architecture that exists in the common domain among the CCEB nations. National concerns to meet this proposed architecture would be addressed within each nation’s internal architecture planning.

TEXT CONVENTIONS

106. Italics are used in examples and in text which must be present in exactly the form shown.

CHAPTER 2

ASSUMPTIONS

201. The proposed Griffin architecture is based on some fundamental assumptions on the environment for the Griffin DS, and future requirements.

DATA OWNERSHIP

202. All ownership of data will rest with the sending nation, whilst access to the replicated data will take place within the receiving nation. Nations may convert the contents of received information where necessary to ensure national usability and to support international interoperability (e.g. conversion of addresses to local internal form, addition of default values for attributes not supplied, changes to DNs and DIT structures etc).

LEVEL OF MESSAGING SUPPORT

203. The Griffin DS will support the interchange of information using a messaging transport service which supports attachments. Peer-to-peer replication between nations will initially be implemented by means of LDIF files using the underlying messaging transport service. Other solutions utilizing different data formats may be supported in future.

ACP 133 SUPPORT

204. It is anticipated that a future version of ACP 133 will define a reduced Common Content subset, designed specifically for international interoperability. It is planned that this subset be supported by the Griffin DS. In the meantime, a subset of ACP 133 will be used as defined in Chapter 6, Annex A (Supported Schema). This subset may include proposed ACP 133 schema changes defined in the current draft Edition C.

GRIFFIN INFORMATION DOMAINS

205. The Griffin network infrastructure supports various “eyes” information domains provided between different nations, with “two-eyes” being considered the most secure, and “five-eyes” the least. Nations will transport entries to a given nation over the most restrictive (i.e. least eyes) network connection available between the two nations, thus ensuring that information does not leak to other nations who may not be authorized to access that information. How this is achieved is dependent upon the transport mechanism employed.

DATA SECURITY

206. The CCEB Infosec Working Group have advised that data transiting the Griffin network should be digitally signed to allow the recipient of that data to assure themselves that the data was actually sent by the purported originating person, role or nation, and that the data has not been modified in transit. The use of the MM network will meet this requirement, and hence it is strongly recommended over the use of the Griffin informal email service. However, until such time as the MM service is available between all nations, the informal

email service may still be employed in parallel with the MM transport mechanism. This will allow the directory information necessary to set up the MM service itself to be replicated between nations, although the mechanism for initially sharing addresses to be employed between nations would be subject to bilateral agreement, and may include hand carrying of data.

DATA INTEGRITY

207. There are two primary concerns to be considered in terms of the data integrity of the transferred LDIF file:

a. There is a need to ensure the LDIF file received by a nation has not been inadvertently corrupted. Different transport services provide different levels of capability for ensuring message integrity. In view of this, the Griffin DS will provide its own internal mechanisms to allow data integrity to be assured. This will be accomplished using a Hash value calculated from the data before it is sent and carried in the message to allow its verification at the receiving nation. Note that the Hash is not intended to protect against malicious changing of data since it is assumed that this would not happen on Griffin, and hence no further protection of the Hash is provided.

b. There is also a requirement that sending nations may not include LDIF entries for which they do not hold the master copy. Receiving nations must ensure they reject the entire file where the sending nation does not hold the master copy of all entries received (e.g. any nation receiving a UK entry in the US LDIF file should reject the file, specify the reason for rejection and request a retransmission).

DATA VOLUMES

208. It is anticipated that the Griffin DS will need to support much larger interchange files than was required with previous DS solutions. It has been identified that proposed messaging transport mechanisms may not be capable of transferring these large files unchanged. Additional features which have been added to the protocol to alleviate this problem include:

a. File compression (where allowed by network and security infrastructure). Use of file compression provides not only the possibility of transporting data files which exceed the allowed limit, but is also a more efficient use of the network by reducing network bandwidth requirements (it appears indicate a compression ratio of 10:1 is achievable). Initially, use of ZIP compression is assumed, but this specification does not preclude the use of other methods and products as they become available.

b. File Segmentation. The maximum file size limitation requires that a data file which is larger than a configurable limit is segmented into multiple data files, each smaller than the limit. These will be sent in separate messages, each with a separate file name, allowing the original file to be recreated in the correct order by the receiving nation.

c. Use of a Web Server. The ability to download data files from a web server can not only allow larger files to be transferred than can be supported through the messaging services, but can also offer a more effective data transfer mechanism, reducing bandwidth requirements. A URL specified within a message can be used to access the data file on the specified web server.

d. Offline Transfer using External Media. It should be noted that this mechanism has been added as a contingency and is not intended for general use. It should be used only for initial transfer of information before the underlying MM service is available or for use in transferring large data volumes in circumstances where other mechanisms for handling large data volumes are not available. It may also be useful for operational situations outside of Griffin such as deployments (where available bandwidth is limited) or for data interchanges with mobile platforms or entities which are not permanently online.

SCOPE OF THE GRIFFIN DS

209. This architecture is limited to addressing the current Griffin CCEB coalition. It does not presently address external coalitions (e.g. MIC Nations, NATO etc).

CHAPTER 3

GRIFFIN REQUIREMENTS

301. This Chapter considers further requirements identified by the Griffin Nations in order to meet their international messaging and data sharing obligations and the effect of these on the Griffin DS.

GRIFFIN INFORMATION SHARING REQUIREMENTS

302. The primary motivation for Directory Services is to effectively exchange contact and addressing information with other nations, thus allowing those nations access to messaging addresses for use when composing either Military Messages or informal emails, as well as to contact persons and/or role incumbents in peer nations. In order to achieve this, information shared must be tailored for effective use within the recipient nation, not for ease of creation or management within the sending nation. Thus there is a critical need for information such as telephone numbers to be presented in a way which is useable (e.g. international format numbers, not local numbers).

303. Message addressing is even more critical. Different addressing attributes are used for the different networks, X.400 O/R Addresses in the case of MM and SMTP Addresses in the case of email. Previous Griffin Directory Services either only supported a small number of entries and attributes (Griffin Initial DS) or only supported informal email (Griffin Interim DS). The Griffin DS is intended to handle both informal email and MM. The following sub-paragraphs discuss the specific requirements of the two messaging infrastructures.

304. The task of the Griffin DS is not only to allow exchange of information between international users who have direct access to the Griffin network, but also to ensure that it supports any additional needs which may not at first be apparent. One particular case has already been identified and this is documented in paragraph 317. This does not preclude further requirements and exceptions being identified and documented in future.

MILITARY MESSAGING ADDRESS REQUIREMENTS

305. It is understood that only a single X.400 Military Messaging address will be required for sharing through the Griffin DS for each Organization, Address List, Person or Role which a Nation chooses to share. If this external address differs internally from internationally, it is the responsibility of the sending nation to only expose the international address, and to perform address translation within its national domain. This implies that the complexities of addressing within information domains required by the informal email service, or the need for different addresses dependent upon the security label applied to the message are not required.

306. The Griffin DS assumes the use of the *mhsORAddresses* attribute for storage and transfer of the X.400 Military Messaging address. This attribute is defined as multi-valued, and hence more than one address could be transferred within it. However, sending nations need to understand that it is unlikely that receiving nations would be able to differentiate

between the different values or know which would be the correct address to use in different circumstances. The address selected by the received nation may be indeterminate.

INFORMAL MAIL ADDRESS EXCHANGE REQUIREMENTS

307. The *mail* attribute is used to hold informal email addresses, suitable for use over Griffin. In order for replication of directories to support the informal email mode of addressing within a nation, multiple email addresses (each relative to an allowable information domain subset) must be replicated for each addressable directory entry; one for each information domain a particular entry can participate in between the two nations.

308. The Griffin network email addresses are created with a structure as follows:

a. user@xxxxx.yyyyy.informationdomain.ISO3166 2char countrycode

b. The *user*, *xxxxx* and *yyyyy* in the structure above represent portions of the email address that a nation can populate with no internationally imposed restrictions. The information domain and ISO 3166 two letter country code portions must follow internationally agreed standards for format, with the exception of the United Kingdom who will use UK rather than GB as their 2 letter country code in this context. The format of the information domains has been agreed between the participating nations and is specified by the Coalition Wide Area Network (CWAN) Working Group (WG) in Section 1 of Tab 2 of Annex C to CWAN-CCEB-DOC-TECH-02-019-VER1.0.0.

309. Examples:

a. *Howell.am@AUSCANNZUKUS.ca* would be an address for a Canadian user within the five-eyes domain.

b. *smithj@js.pentagon.UKUS.us* would be an address for a US Joint Staff user within a bilateral domain between the UK and US.

310. Since the email address attribute is multi-valued and can contain addresses for multiple information domains, all values may not be sent to all nations without validating each value and removing it, if it is not appropriate to be sent. It would be possible for a nation receiving a value which did not contain its own nation to derive information about domains in which the providing nation is participating. To avoid this scenario, it is the sending nation's responsibility, to send only email addresses that contain at least both source and destination country codes for the information domain that the email is being transmitted over.

311. For example, the following three addresses might be stored in one entry's email attribute:

a. user@js.pentagon.AUSCANNZUKUS.us

b. user@js.pentagon.UKUS.us

c. user@js.pentagon.CANUS.us

312. If a bilateral is established between the US and UK for directory replication of this entry, the *user@js.pentagon.CANUS.us* value in the email attribute may not be shared over this domain because it could compromise the CANUS information domain. If a nation receives more than one value in the address field and stores the attribute unchanged in the directory entry, it should be recognised that the order in which the values are returned to a user reading the entry is not determinant, and could lead to the wrong value being used, possibly compromising security. It is therefore the receiving nation's responsibility to ensure that addresses are presented to their users in a usable and user-friendly manner.

MISSING MESSAGE ADDRESSES

313. Except when an entry is supplied to provide structure to the DIT, if that entry does not contain either an email or X.400 address that corresponds with the replication peer (i.e. Right Hand Side (RHS) of the email address does not contain the replication peer nation) a nation must restrict the entire entry from being replicated[1].

LEGACY MESSAGE ADDRESSING AND ROUTING

314. One of the identified tasks of the X.400 Military Messaging service is to provide a replacement service for the existing legacy messaging solutions between nations, such as ACP 127. It is envisaged that nations may choose to cut all legacy links to nations with whom they have Griffin MM connectivity even before all legacy systems have been removed internally. In this case, X.400 MM messages may be received which need to be forwarded internally to a legacy recipient. In the case of an ACP 127 legacy system, signal addressing is through use of Plain Language Addresses (PLAs) with the originating PLA being passed in the signal as well.

315. There is a requirement for the originator X.500 Distinguished Name to be converted to a PLA during the conversion process to an ACP message. This capability can be facilitated. It is agreed therefore that the *plaNameACP127* attribute (part of Auxiliary Object Class *aCPPlaUser*) be exchanged for this purpose.

316. The use of the Griffin Military Messaging service and the Griffin DS as a conduit into legacy messaging systems (as well as its possible future use in other infrastructures such as to deployments) identifies a future need for the exchange of information regarding the receipt and display capabilities of recipients. At the moment it is assumed that all recipients can handle equally large messages and can handle any attachment included by the sender. Since this assumption is not correct, the need for the exchange of capability data identifying what a recipient can or cannot accept becomes paramount. This capability is not included in the

[1] This policy may require further consideration. The possibility of exchanging entries not containing addresses, but possibly only telephone numbers and other contact details may need to be allowed, but only when a better mechanism for exchanging allowable domain releasing information has been specified. Currently the SMTP address is required to identify which "eyes" domains the addresses contain and hence may be used to replicate the entry into the appropriate domain(s). In the absence of such information, it would not be possible to handle the entry correctly.

Griffin DS at this time, but this paragraph has been included as a placeholder for its future incorporation, in parallel with it being defined within ACP 133.

GRIFFIN MESSAGING SERVICE REQUIREMENTS

INTRODUCTION

317. The Griffin DS itself also requires the use of an underlying transport service over which directory information to be replicated may be transferred. Both underlying transport services, the High Grade, MM service and the informal email service could be used. The preference is to use the MM service, when available, thus ensuring information is digitally signed and better protected. The exchange of addressing information is heavily based on the architectures adopted by nations to implement their internal and boundary solutions, as discussed below.

MILITARY MESSAGING SERVICE

318. The Griffin MM Service is based on the transport of X.400 messages across the Griffin network via ACP 145 messaging gateways provided at the boundary of each nation. How the service is provided behind the ACP 145 gateway is a national matter, as is the choice of gateway within a nation (if more than one exists) and the choice of security domain used to deliver the message (as long as it does not contravene security rules).

319. The Griffin DS requires that the MM service supports the following capabilities:

a. Transfer of text attachments.

b. Ability for the originator to specify security labelling information in both the message heading and the First Line of Text of the Message text area, to allow the Classification and Releasability requirements for the message to be specified.

c. Messages signed at the gateway on behalf of the nation.

d. Choice of an appropriate path dependent upon the security classification which was assigned to the message before it was sent.

e. Message delivery to the specified destination mailbox.

INFORMAL EMAIL SERVICE

320. The Griffin email system is based on an RFC822 SMTP mail service which must work across multiple information domains. The Griffin network currently supports two-eyes information domains between certain nations, in addition to the five-eyes AUSCANNZUKUS domain, although other domains could be added in the future, and are not precluded by the Griffin DS. The most restrictive information domain (i.e. least releasability) is a two-eyes bilateral domain, such as *CANUS, UKUS*, etc. The least restrictive information domain is the five-eyes domain.

321. The need to correctly support and manage messaging over the different information domains results in the following messaging requirements of the informal email service:

a. Text attachments may be attached to messages.

b. Messages which are correctly addressed will traverse only the requested information domain (or one more restrictive).

c. An agreed mechanism is available (such as First Line of Text (FLOT)) to allow the Classification and Releasability requirements for the message to be specified.

d. The boundary guards within a nation will ensure that the releasability rules specified within the message would not be violated if the message were to take the path specified by the recipient address.

322. For messages in transit beyond a national boundary, the RHS of an email address will indicate within which information domain the message must be transmitted, thus allowing the national boundary services to route the message correctly. This may be achieved within a nation by either requiring the originator of a message to specify the correct address from a list of alternatives or by the system automatically putting the address appropriate to the releasability information provided at the boundary between the nation and the common network. If the manual approach is adopted, the need for checking of releasability information against addresses in the message at the boundary is paramount.

DIRECTORY INFORMATION SCHEMA REQUIREMENTS

323. The ACP 133(C) contains an International Common Content, which will define those Object Classes and Attributes which must be supported for international interoperability, as well as specifying syntax rules on how the information should be presented for international use (e.g. format of phone numbers). This will have the benefit of allowing nations to procure products which will support and interoperate with a defined subset of objects and attributes, as well as ensuring that data which is shared internationally is also usable within each target environment.

324. It is intended that directory information replicated using the Griffin DS is compatible with this proposed International Common Content. This is in preference to an artificially defined subset, as was the case with the earlier versions of the Griffin DS. In order to ensure

that the correct version of the schema has been used for directory data being shared over Griffin, a schema version number will be specified in the control information passed between nations. If this option is not used it is assumed that the nations have agreed upon a common schema through external means.

325. Since it may take some time to get final approval for ACP 133(C) for international use Chapter 6 Annex A defines a subset schema which should be supported by early implementations of the Griffin DS to allow the development and use of MM as well as the currently used informal email transport mechanisms.

326. This subset will include the object classes supported by the Griffin Interim DS:

a. Organizational Unit.

b. Organizational Person.

c. Organizational Role.

327. These object definitions will be extended to support X.400 addressing and other auxiliary object classes and attributes required to support the MM service, such as O/R Address and *plaUserACP127*.

328. In addition, new Object Classes are required including:

a. Address List (used to send military messages to a predefined list of message recipients).

LARGE DATA VOLUME REQUIREMENTS

329. It was stated previously that much larger data volumes are envisaged by the nations for use with the Griffin DS. One nation has estimated potential data up to 100Mb (possible worst case). Current limits on the X.400 MM service are assumed to be a maximum message size of between 2 and 8Mb, and even if this were increased, it is unlikely that a messaging service is the best choice for data volumes this large:

a. It is not clear that these sizes are necessarily accurate, rather being a worst case estimate. Even so, it is accepted that volumes are likely to be much larger than a 2Mb limit.

b. Only one nation is likely to send volumes of this magnitude. Other nation's data will be much less, although it is not certain that existing maximum message sizes would necessarily be adequate for other nations.

c. A full transfer is only required initially and when synchronisation is lost. Otherwise, data transfers are limited to changes only, which are likely to result in much smaller files. A solution which could handle large volumes poorly but smaller volumes more efficiently may be acceptable.

330. A number of alternative solutions have been considered for resolving this issue, including:

a. Attachment Segmentation. One solution to this problem is to split the data to be transferred into acceptably sized chunks, and to transmit these in successive messages, reconstituting the data after all has been received. Whilst this is relatively easy to implement from a sending perspective, numerous difficulties in error situations make it less attractive for the receiver. Error recovery may need to be handled manually rather than through an automated mechanism.

b. Data Compression. It is unclear whether compressed data would be allowed through the boundary guards of each nation. Even if it is, the choice of compression algorithms may be limited. In addition, initial analysis has indicated a compression factor of 10:1, which would still cause data to exceed the messaging limitations in some cases. Since compression has an obvious network capacity benefit where it is allowed, the optional use of compression has been added to the Griffin DS protocol, even though in some cases its use may not actually be possible. Currently only support for ZIP has been included although other algorithms could be added in future.

c. Use of Subtrees. The Griffin DS protocol supports the possibility of transferring data as multiple, separate subtrees. If the sending nation's data is structured appropriately, this can provide an alternative solution to segmentation.

d. Use of Web services for accessing data. The fourth solution requires the use of web technology. Rather than sending an attachment containing the data, a URL reference could be sent instead, from which the data may be downloaded, using FTP or other allowed methods. Web access in different Information Domains across Griffin has not yet been agreed upon and a complete service is unlikely before the inception of the Griffin DS.

e. However, it appears possible that nations could unilaterally set up web servers in the Griffin network space for generalized or specific access. Since it is possible that only one nation will actually have large data volumes, the creation of a single web server for them to post information for each peer nation may be an attractive solution. The exact logistics, security implications and information exchange rules have yet to be worked out. However, the solution appears to be sufficiently general and workable to include the specification of a URL as an alternative to a File Name (and associated attachment) within the protocol.

f. Offline Data Transfer. Another solution which has been added for contingency purposes is support for the offline transfer of data through external media. It is not intended for general use but should be supported by all Nations for reception of data.

DATA FORMAT REQUIREMENTS

331. LDIF is currently the data format of choice for use on the Griffin network for transporting directory information. Most directory server products, both X.500 and LDAP servers, can support LDAP and in many cases provide built in tools for creation and consumption of LDIF data. If not, there are numerous COTS products which can process LDIF files, converting to LDAP for directory access. It is however recognized that this is an area of rapidly changing technology and that in future, other mechanisms may be more appropriate.

332. Without changing the underlying file based mechanisms on which the Griffin DS is predicated, it is possible to envisage alternative data formats which could be employed in the future for reasons of efficiency or compatibility. A possible candidate alternative is use of DSML, an XML variant designed for directory access, although other alternatives may be equally suitable.

333. It is not planned to modify the Griffin DS protocols to support different formats at this time, although a capability has been incorporated which can, in future be used to support alternative formats.

CHAPTER 4

GENERAL ARCHITECTURE

OVERVIEW

401. The Griffin DS architecture requires that directory information be held by each participating CCEB nation behind their national border and shared via data files, with LDIF being the first replication file format supported. This file based directory replication will be peer-to-peer, will include addresses for use on all information domains allowed between the replicating nations, and will be sent over the most restrictive information domain available between the two nations (see Figure 4-1).

402. A common schema will be supported based on the International Common Content subset of ACP133(C). Until this subset is ratified, the supported schema is defined in Chapter 6, Annex A. Data Integrity mechanisms implemented will provide some degree of additional protection over the security provided by the Griffin network. Policies and operational procedures are discussed along with national responsibilities in this paper.

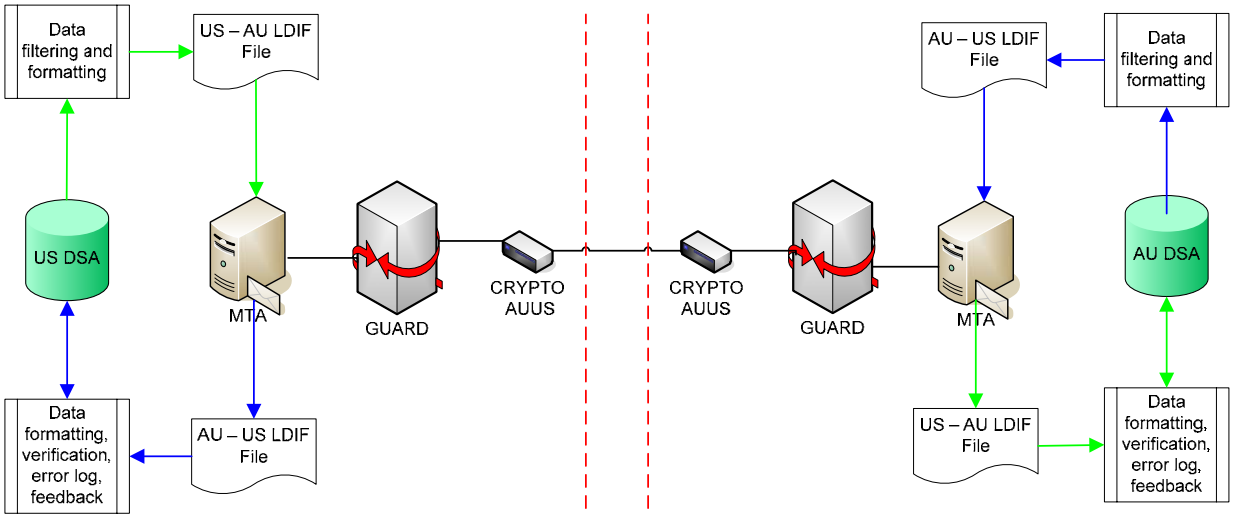

Figure 4-1 Peer to Peer Australian/US Replication

TRANSPORT MECHANISMS

OVERVIEW

403. An agreed mechanism is required to transfer the data file and associated control information between national directory managers. This will initially be based on a messaging service, but may, in future also support alternative mechanisms such as the use of FTP or by

posting information onto the Web, as and when these services are supported over Griffin[2]. LDIF files will initially be sent as text attachments to messages using either the MM service or the informal email service, as agreed between nations on a per replicated area basis.

404. A nation's directory information will be converted into a file and attached to an SMTP or Military Message for distribution to other nations of interest on a peer-to-peer basis. Each nation is responsible for filtering entries and values from their national directory to ensure the attached data contains information releasable to the receiving nation. This provides a means to identify corrupted file segments permitting retransmission of the bad file segments rather than resending the entire LDIF file.

REPLICATED DATA VOLUMES

405. It has already been noted that data volumes to be transferred may exceed the maximum allowed by the messaging service for transport within a single message. File segmentation has been added as an optional capability, whereby a large file is segmented into separate parts, each of which is sent within a separate message. This mechanism extends the existing file sequence numbering scheme to incorporate a segment sequence numbering scheme to allow all parts of a single file to be received in the correct order and processed correctly. Individual checksums for each separate part of the file sent augment the complete file checksum, which is sent with the first message in the sequence. This provides a means to identify corrupted file segments permitting retransmission of the bad files segments versus resending the entire LDIF file.

406. In addition, further support has been added to alleviate the problem by use of separate replicated areas, compression or posting files to websites, although these mechanisms will be dependent on individual national capability or security regimes. The file segmentation mechanism is therefore currently recommended for most situations.

SPECIAL DS MANAGEMENT MAILBOXES

407. Nations should create special SMTP and/or X.400 mailboxes that will be used specifically for sending and receiving LDIF files. A number of guidelines have been developed relating to these:

a. Nations may create a single account for both purposes or may create separate accounts to send and receive replicated LDIF files. It is recommended that the left-hand side (LHS) of an SMTP email address for receiving/sending email should be *diradmin* as should the *CommonName* attribute of an X.400 address.

b. Nations may also wish to establish separate accounts per bilateral/multilateral agreement.

[2] It should be noted that changes to use other than messaging for data transfer (or data location at a web site) would represent a major change to the protocols and would require a major rewrite of this specification.

c. Nations must communicate these accounts, and their purpose, to their replication partners.

SUPPORTED SCHEMA AND DIT

SCHEMA

408. The Griffin DS will be based on the core Common Content schema from ACP 133(C). Until this is available, a working subset has been defined in Chapter 6, Annex A, which supports the following entry types:

a. Country (only allowed for High-Level DIT).

b. Organization (only allowed for High-Level DIT).

c. Locality.

d. Organizational Unit.

e. Organizational Person.

f. Organizational Role.

g. Address List.

409. Griffin contact and addressing information will be mainly populated based on units, persons and roles. The following is a list of some of the more useful attributes:

a. Entry name (e.g. common name) of Unit[3], Person or Role entry.

b. Personal details (e.g. Surname, Given Name, Title, Rank etc as appropriate).

c. Pointers between role and person entry (where appropriate).

d. Phone numbers (e.g. Phone, Fax, Mobile IP etc). These should be provided in international format (e.g. +1 7031234567 with no spaces other than between the National Dialling Code and the number).

e. X.400 MM address.

f. Plain Language Address (PLA) for legacy message conversion.

g. SMTP email address which is multi-valued and should contain values for each “eyes” information domain through which the entry can be addressed.

[3] OU attribute for Organization Unit.

DIT STRUCTURES

410. The high-level DIT for each nation will be static and is currently defined in Annex B. It is the responsibility of each nation to create all other nations' high level DIT structures in their local directory.

411. Nations may wish to replicate more than one subtree of the DIT to other nations (for example, role, persons and PKI information may be held in different areas of the DIT). The sending nation must replicate the lower-levels of their DIT, up to the agreed top level DIT entry to each other nation for each separate area replicated. To assure data integrity, nations should wherever possible maintain the replicated DIT structure defined by the sending nation.

DIRECTORY DATA REPLICATION MESSAGE FORMAT

412. The Griffin DS protocols are based on the transfer of Control Information along with the data file containing the directory information. This control information is described in Chapter 5 of this document. A few general guidelines are contained in this chapter to aid understanding of this control information.

413. All forms of replication will consist of a data file (typically as an attachment of the message) which conveys replicated data from the sending nation and a message body text used to transport the Control Information about the replicated information to the receiving nation.

414. The Griffin DS Control Information will be contained primarily in the Message Body Text area of the message. In addition, the name of the data file will be contained in the Message Subject field, and where appropriate, much of the Control Information is duplicated at the start of the data file. In the case of an LDIF file, this information is present as comment lines, except for the Hash value which can only be calculated after the data file has been finalized. When other data formats are supported, if they allow the use of comment lines, it is anticipated that the Control Information would be similarly duplicated within them. This duplication is primarily present to allow a file, which has become separated from its associated message, to be re-correlated.

DATA SECURITY

REPLICATED DATA INTEGRITY

415. Each data file must have checks performed on it by both the sending and receiving nations to maintain DIT integrity and prevent security compromise of the information domains. It is each nation's responsibility to assure that the replicated information contains only data releasable to the receiving nation.

416. To assure the integrity of the data file, the sending nation will send a hash of the file within the Control Information. This hash will not be digitally signed and therefore does not require PKI.

417. Nations will normally only pass changes since the last update in the data file. To assure synchronization of the replicated DIT structure with the mastering nation, nations should send a complete refresh (full national DIT structure excluding high level DIT) to the replicating partners on an agreed schedule (see paragraph 462).

418. SHA-1 shall be used for the data file hash. SHA-1 produces a 20-octet message digest from a source, in this case the data file. This digest will be relayed to the replication peer as a text-based representation of the hexadecimal digest value (i.e. 40 ASCII characters that represent digits 0 - 9 and characters A – F).

DATA CONFIDENTIALITY

419. The underlying network on which the email and military message is being transported will supply confidentiality of replicated data. There is no requirement for encryption beyond the underlying network.

DATA ACCESS CONTROL

420. It is assumed that all shared national DIT entries would be for unrestricted read access. Thus, there is no need for nations to include access control information in the data files.

421. Receiving nations must ensure that users cannot modify the replicated DIT structure. How a nation protects replicated DIT entries from unauthorized modification is a national matter and outside the scope of this document.

RECEIVED DATA CONTENT INTEGRITY

422. As well as the presence of a hash checksum to ensure that the information has traversed the network without inadvertent changes, nations must ensure that the data file received contains valid information from the sending nation. Checking could include:

a. Validation that the file does not contain entries associated with another nation's DIT structure (e.g. UK entries in a US LDIF file). This may be achieved by checking the distinguished names (DNs) of all the entries in the LDIF file to verify correct placement of those entries in the sending nation's DIT.

b. Validation that the country code of the sending nation is contained within the RHS of all email addresses within the LDIF files they are receiving. It is unlikely that a similar check is required for X.400 addresses although this will depend on the messaging architecture.

SECURITY LABELING

MESSAGE SECURITY LABELS

423. The SMTP protocol does not support the concept of Security Labels within messages, other than conventional and mutually agreed use of the Subject field or First Line of Text (FLOT) to hold a label. No further use of Security Labels is considered here, although most nations will impose rules on how messages are to be formatted to ensure they can be releasability checked at the national boundary. The Griffin DS has specified a conventional format within the first line of the message text which is suitable for most nations (see paragraph 509).

424. The X.400 Military Messaging service requires a Security Label to be inserted in the S/MIME v3 ESS SecurityLabel attribute by the ACP 145 Gateway, and it is assumed that this will be derived from a Security Label specified implicitly or explicitly when the military message was created. Any Security Label so generated must reflect the classification and releasability of the message, especially if the gateway uses this information for message routing. The conventional first line of message text should still be present for military messages.

SECURITY LABELLING OF THE MESSAGE CONTENT

425. The message security label (for X.400 messages), message body text and the data file (where appropriate) should all contain security label information. Security label information may include policy, classification and releasability information. The policy identifies the national policy against which the rest of the labelling information applies. The classification level of the message and data file are determined nationally, whilst the releasability reflects the information domain over which the message should be transmitted and must be set to the most secure domain implied by the contents of the data file information.

426. The syntax of the policy, classification and releasability information is specified nationally (generally in a user concept of operations document) and may need to be included in the first line of the message text (FLOT), as well as in the Security Label for military messages. The receiving nation should respect the classification and releasability indicated by the sending nation and store/process the data appropriately. As well as specifying security labelling information in the FLOT and/or message label, the Griffin DS service requires the originating nation's security label to be present in both the Control Information in the message text and the data file.

427. The general notation for Griffin DS labelling information is a text line of the form:

Classification policy prefix classification releasability

428. This should appear in the message text and the data file (where appropriate). However, it has been identified that certain systems may be unable to present the first line of text in this form due to internal implementation restrictions. Some flexibility may be required in the exact format of the information.

429. In the message body text, it must appear as the first line of Griffin DS Control Information. In the data file, it must also appear as the first line of commented out heading information.

TYPES OF REPLICATION

430. There are three types of replication used for the Griffin DS. These are **initial, incremental** and **verify**, and each is described further in the following sub-paragraphs.

INITIAL REPLICATION

431. This type of replication comprises a full replication of all data to be shared below the high level DIT and is used to initially populate a nation's lower level DIT. This full replication requires a sending nation to include all entries in their low level DIT that apply to the receiving nation. The data will contain all entries, attributes, and values that a nation is willing to share with the receiving nation. In the case of LDIF format data files, variations in production of these files requires support on receipt for the inclusion of **changetype: add** LDAP operations.

432. A full replication may also be needed between nations to solve operational problems. The frequency of these total refreshes will be by agreement between the replication peers. Nations must ensure that a capability is available to generate an appropriate initial replication file on request from a peer nation within a mutually agreeable period of time. In addition, a sending nation may need to reset shared directory information due to some internal operational reason. In this case, a nation may send an Initial Replication at any time with a sequence number of zero. Where practical, the sending nation should contact the receiving nation offline to indicate this occurrence.

INCREMENTAL REPLICATION

433. This type of replication will update the sending nation's lower-level DIT with changes that have occurred from the last data exchange between the replication peers. The data file will contain operations to be performed on the receiving nation's local directory to bring it into alignment with the sending nation's copy of the DIT. Where LDIF format data files are used for incremental updates, changes are expressed using the LDIF **changetype** syntax to add entries, delete entries, modify existing entries and modify an entry's DN.

434. Incremental replication will be performed on an agreed frequency, except for emergency situations where updates need to be sent outside of the normal schedule. Therefore, in order to maintain the replication frequency, if no changes have occurred since the last scheduled incremental replication, a "blank" data file will be sent, containing no

directory entries (i.e. the hash value will be a hash of the Control Information fields (if present) within the data file).

435. Sending nations must ensure when producing this form of replication that they maintain information that relates an entry to the proper information domain. For example, if a nation builds their data file from a directory product's change log, deleted entries may no longer carry the email attribute needed to determine the information domain. Additionally, modify entries in a change log also may not contain needed information to determine the information domain. It is the responsibility of each nation to ensure the accuracy and quality of its outgoing data files.

436. It is advised that receiving nations archive (at least) the last received **initial file (**or verify file if it can be used to initialize the DIT**)** replication, and all subsequent **incremental** replication files, to allow the reconstruction of the sending nation's DIT.

VERIFY REPLICATION

437. This form of replication is similar to an initial replication of a sending nation's lower-level DIT. By agreement, nations will pass a data file containing their complete lower-level DIT. The intent of this exchange is to allow the sending nation's replication partner to verify their local directory is synchronized with the sending nation's master copy. This form of replication is meant to supplement incremental replication, but not replace it.

438. The sending nation shall ensure that all applicable incremental updates have been provided to the receiving nation prior to sending a verify file.

439. The receiving partner may choose to update their local directory with this complete update, or may choose to compare their local directory DIT against this file either for verification purposes or to correct any differences. How a nation assures synchronization is a national responsibility; however, it is each nation's responsibility to ensure they accurately reflect the sending nation's DIT.

440. If a receiving nation will be performing the compare and verify method, they must apply all outstanding **incremental** data files before using this form of replication. Nations must also take care to recognize the differences, including deletion of entries that exist in the local directory where they are not included in the **verify** data file and addition of entries when the reverse is true.

441. If a receiving nation will be completely replacing their local directory with a sending nations lower-level DIT, any outstanding **incremental** replication data files may be discarded. To ensure the successful update of all entries, it is recommended that nations completely remove all lower level DIT entries prior to applying this form of data file. It is also recommended that the verify data file is tested offline before being used on a live service and that facilities are provided to revert to the last correct version in the event of failure.

INDICATION OF TYPE OF REPLICATION

442. The type of replication being requested will be specified within the Control Information in the message body text, as follows:

a. An Initial replication type should be expressed as *type:initial* and in the file name as *init*.

b. An Incremental replication type should be expressed as *type:incremental* and in the file name as *inc*.

c. A Verify replication type should be expressed as *type:verify* and in the file name as *ver*.

SCHEMA VERSIONS

443. Successful exchange of directory information over Griffin requires that common schemas are used by the sending and receiving nation. In order to achieve this, a Schema Version number may be included in the Control Information. The version number will comprise two or more characters, an alphabetic followed by one or more numeric characters. The alphabetic defines the Edition of ACP 133 in which the schema is defined, and the number defines the schema point release in use. Use of the subset schema defined in Chapter 6, Annex A of this specification is given a generic schema version number of C0. The first ratified version of ACP 133(C) will be given the schema version number C1, with future point releases increasing sequentially.

REPLICATED AREAS

444. A nation may choose to replicate more than one subtree of its DIT to one or more nations. In this case, an optional Replication Area Number may be specified in the Control Information. Separate replication areas are useful if the amount of data to be replicated is large, or the schedule against which updates are required differs markedly (e.g. if PKI revocation information is to be replicated hourly (or less), whilst other information is to be replicated weekly).

445. Replicated areas are identified by the keyword *area* within the Control Information followed by a number to uniquely identify it. If no replicated area is specified, a default area of *0* will be assumed (which would indicate a legacy Interim file). Within a replicated area, sequence numbering of replications will be separately managed. A nation must agree with another nation that it wishes to use more than one replication area, and also specify the High-Level DIT entry below which each replicated area will manage entries. Limited overlap of higher level DIT entries (i.e. Organizational Units) is allowed between replication areas, but this is to be discouraged. Overlapping of data entries and large areas of the DIT is not normally allowed, and a receiving nation may discard duplicate entries, which may in turn lead to verification problems.

SEQUENCING OF GRIFFIN EXCHANGES

446. Since a receiving nation cannot fully verify the correct sequencing of received Griffin exchanges based on creation date alone, Datafile Sequence Numbering is used to ensure continuity of import operations and to mitigate problems that could cause potential import errors or corrupt data because a receiving nation unknowingly missed a Griffin exchange (e.g. system error causing one or more messages containing data file attachments not to be sent or received, or even duplicated). Each sending nation will maintain separate Datafile Sequence Numbers (DSNs) starting at zero for each receiving nation (or Replication Area per Nation), and will increment the corresponding DSN by one each time a Griffin exchange is sent to a receiving nation. Once a DSN reaches 9999, it should be reset to start at zero again.

447. It is the responsibility of a receiving nation to keep track of DSNs from sending nations such that it can be determined that a data file was not received (e.g. nation last received a file with DSN 74 and now has received a file with DSN 76) or was duplicated. If an exchange is missing, the receiving nation should email or send a military message to the sending nation requesting the missing exchange, specifying the particular DSN or requesting a new initial or verify data file (as defined in paragraph 208). If an initial exchange is sent, it will contain a new DSN of zero. Hence, it is the responsibility of the sending nation to archive Griffin data files and have the capability to process such special updates.

448. It should be noted that where segmentation is in place (see paragraph 415), the same DSN will be used for all segments within the exchange.

DATA FILE SEGMENTATION

449. In the case of large data sets, which exceed the maximum message transfer file limits, and where compression or other mechanisms are not suitable for reducing the size of the data, the use of Data File Segmentation may be considered. It is strongly advised that the Maximum Data File size which may be transferred is configurable, as should be the Maximum Segment Size which may be sent. Where segmentation is required, the data set is segmented into multiple files (each smaller than the configurable Maximum Segment Size), each of which is sent in order in a separate message.

450. Management of the receipt of the segment and the re-composition of the original data set is the responsibility of the receiving nation, which may need to request retransmissions from the sender nation if required (see paragraph 465). In order to assist with this process, the message header is modified in a number of areas, maintaining backwards compatibility as much as possible.

451. The file naming policy is modified to support segmentation by placing a file segment number following the Universal Time (UTC), with zero being used to indicate a non-segmented file. The Sequence Number line is modified to add additional (colon separated) fields where sequencing is being performed. Following the Datafile Sequence Number, which is the same for all messages in the sequence, is a Segment Sequence Number (starting at 1 for

the first segment), a Highest Segment Sequence Number and a Hash Value for the complete Data set (the latter two being mandatory for the first segment, optional for others).

452. Segmented data files (other than the first) will not contain a heading since this would make it impossible to reconstitute the file exactly on receipt, and thus would invalidate the complete file checksum.

453. Experience with certain Griffin gateways has indicated that additional End of Line (i.e. CR/LF combinations) may be added unless a text attachment terminates in an End of Line sequence, which causes checksum checking to fail. In light of these discoveries, it is required that segmenting only occurs at the end of a line when sending text files. Implementations should preferably truncate files on the line before the segment size if reached, although receiving nations should cope with slightly more. Implementations should also ensure that segment sizes slightly larger than the agreed and configured limit should be supported on receipt to cater for the case where segmentation occurs at the line immediately following the maximum segmentation size (since LDIF lines are typically 80 characters or less, an increased allowance up to say 100 characters should be adequate).

454. It should be noted that while segmented messages should be sent in ascending sequence, they may not be received in the same order, since the underlying messaging systems cannot guarantee order of delivery. The precedence of messages is for bilateral agreement between nations, but may be dependent upon message size, urgency etc.

DATA FILE FORMATS

455. Replicated data will initially be sent using the LDIF data file format [Ref. 3]. However, it is envisaged that in the future, other data file formats may be allowed, although these have not yet been defined. It is appropriate therefore to incorporate a mechanism to allow the format of a data file to be specified by the sending nation for use by the receiving nation. It is the responsibility of the receiving nation to accept all supported formats.

456. The format of data file transmitted may be specified by the optional keyword *format* within the Control Information in the message body text, followed by a keyword specifying the format of the data file. The only value currently allowable is *LDIF* (which may be specified). If the keyword is not present, the default of *LDIF* is assumed.

COMPRESSION FORMATS

457. The use of compressed data, is allowed by some nations, dependent upon security and networking constraints. Compressed data would greatly reduce the amount of network traffic in transferring large data files. An optional keyword c*ompression* is allowed in the Control Information to specify that compression is being used. Allowable values are currently restricted to *NONE* indicating no compression is present or *ZIP* indicating use of Zip file compression. If the keyword is not present, the default of *NONE* is assumed.

458. Where data is compressed, the compression should be performed on the complete set of data before segmentation. Where segmentation is also required, each segment comprises a

chunk of the compressed data, whose size is less than or equal to the maximum segment size allowed. In addition, the rule on segmentation at End of Line is relaxed (since compressed files would typically not contain End of Lines and it is assumed that any implementation which can support compressed file formats will not attempt to add additional characters into the data stream). The file checksum is taken over the complete, data before compression, thus ensuring that it has been uncompressed correctly. Where segmentation is also necessary, individual segment checksums are calculated against the segment of compressed data, to ensure that each is transmitted correctly.

PROCEDURES

REPLICATION POLICY

459. National agreements will be required to specify the frequency and type of replications for each replicated area. There is no requirement that all replications from one nation to others are necessarily at the same frequency or schedule.

460. Replication strategy and frequency should be dependent upon their data replication requirements and capabilities. A recommended starting point is weekly for incremental and monthly for verify replications. However, nations must allow for the possibility that emergency replication requests may be required outside of the agreed schedule and at short notice. Manual procedures (e.g. telephone call to the directory administration team) may also need to be employed to inform the partner nation that an emergency transfer has been (or will shortly) be sent.

BACKUP POLICIES

461. It is strongly recommended that directories are backed up before received data is loaded into them, and that a test load is performed on an offline copy of the directory.

CONFIGURABLE OPTIONS

462. The following features should be configurable within a national implementation on a bilateral basis with each other nation);

a. Frequencies for Incremental and Verify replications required per replicated area (default = weekly for incremental and monthly for verify).

b. Maximum Data File Size (default = 2,000,000 bytes). This value indicates the maximum files size which may be sent before segmentation should be employed. This value and the next must be agreed between peer nations.

c. Maximum Segment Size (default = 1,500,000 bytes). This value indicates the maximum segment size to be used. The value must be smaller or equal to the Maximum Data File Size.

d. Precedence's used for different types of message (default = *priority* for incremental, *routine* for initial and verify and *immediate*, for emergency updates).

e. String used to indicate start of Griffin DS Control Information (default = "$&#$&#").

f. String used to indicate end of Griffin DS Control Information (default = #&$#&$).

g. Timeouts required for waiting for the receipt of all parts of a segmented message after the first has been received (default = 24 hours).

463. Implementations should also be capable of:

a. Resetting Datafile Sequence Numbers.

ERROR RECOVERY

464. The following error recovery situations must be handled by a National Implementation:

a. In the event of data errors or inconsistencies in the received information (e.g. checksum failure, entries not owned by sending nation, incorrect country codes in addresses etc), the whole file should be rejected and the sending nation should be informed so that they can correct and resend the data file with the same Datafile Sequence Number.

b. In the event that loading the data into the receiving Nation's directory results in errors occurring (e.g. Deleted entry does not exist, Added entry already exists etc), the nation should verify that all local procedures have been correctly performed. If the error is believed to be in the received data, the nation may choose to reject the message and ask for investigation and retransmission by the sending nation, or it may choose to accept and continue with the data. In the latter case, the sending nation should informed, and possibly a Verify operation requested.

c. In the event of Datafile Sequence Numbers being received out of sequence and messages are identified as missing, all later files may either be rejected or retained and the sending nation should be informed so that they can send any missing data files, before (if necessary) resending subsequent data files. If a receiving nation chooses to retain all subsequent files, it may process them following receipt and successful processing of a missing file.

d. In the event of Segment Sequence Numbers being received out of order, the data file should not be processed until all segments have been received. In the event of all segments not being received within an agreed period, the sending nation should be requested to resend any missing segments. Any duplicate segments should be ignored.

e. In the event of the constituent data files of a segmented sequence not being restored correctly or the hash of the complete file being incorrect, the whole sequence of files should be rejected and the sending nation should be informed and requested to send the complete sequence of messages again, using the same Datafile Sequence Number.

NATIONAL RESPONSIBILITIES

465. The following National Responsibilities have been identified:

a. Nations will determine which entries and values they are willing to share with the peer nation. They are responsible for filtering entries from their national directory system.

b. Nations are responsible for applying LDIF and other data format files sent by other nations to their national directories. LDIF files will conform to RFC 2849 [Ref. 3].

c. All access to the shared directory will take place within a given nation. Nations may, by agreement, allow access to their directory system to other nations. However, the definition of such an arrangement is outside the scope of this document. It should be further noted that where a nation chooses to share another nation's information with a third nation, it is the sharing nation's responsibility to ensure that no leakage of classified information can occur.

d. Nations are responsible for defining the means by which the replicated data is accessible from the national email and MM systems (e.g. web interface, loaded into GAL, etc).

e. It is incumbent on the originating nation to build email addresses that correspond to the allowed information domains for each user. For example:

(a) 5-Eye Canadian: doej@xxx.yyy.AUSCANNZUKUS.ca

(b) Canadian CA - US bilateral: *doej@xxx.yyy.CANUS.ca*

f. Only Griffin informal messaging addresses will be populated by the nation (no national internal SMTP based email address should be populated within the CCEB). Likewise, only externally usable X.400 addresses should be populated.

g. Nations will need to derive the email values to be replicated within an information domain from the RHS of email addresses.

h. Nations must ensure that any changes associated with the replication partner are chronological and not included multiple times. LDIF files will be imported in a timely fashion and in the correct order that they are received.

i. Nations will individually backup the data sets that they send and receive to/from each replication partner on a time schedule. Additionally, nations will backup access, audit, and error logs in accordance with national security policy.

CHAPTER 5

MESSAGE FORMATS

501. This chapter defines the formats for messages sent using the Griffin DS protocols. It covers:

a. File Names.

b. Message Subject Field.

c. Message Body Text.

d. File Headers.

502. Finally some further clarification notes and an example are given.

DATA FILE NAMES

503. The format of data file names is defined here. The name includes the sending and receiving nation's mnemonics, the UTC time of file creation, the SSN, the type of replication and the security label. The file type will be dependent upon the type of data contained within it. For LDIF files, the default file type will be *.txt*. For files compressed with ZIP, the default file type will be *zip*.

504. The security label information is placed at the end of the filename so as not to affect the sorting of the LDIF files, placed in parentheses for readability, and normalized with spaces replaced by Underscores ("_").

505. LDIF attachment files names will be created using the following rules:

a. ISO 3166 two letter country code of sending nation (except GB which will be UK).

b. ISO 3166 two letter country code of the intended receiving nation (except GB which will be UK).

c. Zulu (UTC) creation date and time of the LDIF file creation as follows: *YYYYMMDDhhmmssZ*.

d. Underscore ("_") followed by the Segment Sequence Number (SSN) (defaulted to 0 if the file is not part of a segmented file)

e. Underscore ("_") followed by replication type, represented by *init* for initial, *inc* for incremental, and *ver* for verify.

f. Underscore ("_") followed by the normalized security labelling information to be appended by Nation in parentheses. Normalization in this case is the replacement of spaces by underscore characters.

g. An extension consisting of a period followed by *txt* (*.txt*)

506. Examples of LDIF filenames could be:

a. aunz20030325132559Z_0_init.txt (file is not segmented)

b. caus20031103175442Z_3_inc_UNCLASSIFIED.txt (third file of a segmented file).

507. The filename is included in the Control Information in the message body text and data file header to show the original filename in the event that a messaging system changes the original filename prior to delivery, thus allowing it to be easily restored.

MESSAGE SUBJECT FIELD

508. The email and military message subject field will contain the text:

a. *Griffin DS data:<space>* followed by the data file name (as defined in paragraph 505).

MESSAGE BODY TEXT

509. Due to the uncertainties with the variable format and number of lines comprising the First Line of Text (FLOT) since this could change as the message traverses different national systems, it is not possible to make assumptions of where the Griffin DS Control Information will begin. A configurable character sequence (at the start of a new line) will be used to indicate the start of the Control Information.

510. The message body text will contain the Griffin DS Control Information and will adhere to the following format and rules:

a. $&#$&# to identify the start of the Griffin DS control information.

b. The first line will contain the text *classification:<space>* followed by the classification and releasability information.

c. The second line will contain either:

(a) The text *filename:<space>* followed by the LDIF file name or

(b) The text *url:<space>* followed by the Universal Resource Locator (URL) of the location on the web site from which the file may be downloaded or

(c) The text *offline:*<*space*> followed by the LDIF file name in format of a data file which will be provided to the receiving nation via an offline, out of band mechanism (e.g. hand carried CD-ROM).

d. The third line will contain the text *type:*<*space*> followed by the replication type (*initial* for initial, *incremental* for incremental, or *verify* for verify).

e. The fourth line will contain the text *nation:*<*space*> followed by the ISO 3166 two letter country code of the sending nation.

f. The (optional) fifth line will contain the text *area:*<*space*> followed by a number (0-99) as specified in paragraph 445. If the line is absent, a replication area number of *0* is assumed.

g. The sixth line will contain the text *sequencenumber:*<*space*> followed by up to three numbers (0-9999) and a 40 ASCII character hexadecimal hash value as described in paragraph 451, of which only the first must be present unless segmentation is being performed, each separated by :, as follows:

(a) Datafile Sequence Number (DSN).

(b) Segment Sequence Number (SSN).

(c) Highest Segment Sequence Number.

(d) Hash checksum for the complete (unsegmented) file.

h. The (optional) seventh line will contain the text *schemaversion:*<*space*> followed by a character sting containing a leading alphabetic character followed by one or more digits. These comprise the ACP 133 Edition character followed by the point release number.

i. The (optional) eighth line will contain the text # *format:*<*space*> followed by the type of data file (*LDIF* for the default LDIF format). If the line is absent, *LDIF* is assumed.

j. The (optional) ninth line will contain the text *compression:*<*space*> followed by the type of compression used (if any). *none* and ZIP are the only values initially supported. If the line is absent, no compression (*none*) is assumed.

k. The tenth line will contain the word *hash*:<*space*> followed by a 40 ASCII character hexadecimal representation of the SHA-1 hash derived from the attached LDIF file. In the case of file segmentation, this hash is for the file segment attached to this message.

l. The eleventh line will contain a separator consisting of the string #&$#&$, to indicate that no information beyond this point should be read by systems

automatically processing the message. In this way, any nations that automatically add additional security markings or other information to outgoing or incoming emails can do so without affecting the Griffin DS Control Information.

511. In the case where optional lines are not included, a blank line should not be inserted.

DATA FILE HEADERS

512. Where allowed by the data format used, the first lines of the Message Body Text up to but not including the *hash* line, should appear at the start of the data file to allow the two to be correlated if they become separated, except for the minor differences as noted in paragraph 514. The information will be held in the data file as comments, in the case of LDIF files consisting of a Hash character (# = 0x23) followed by a space (0x20) character in the following format:

a. The first line will contain the text *# classification:<space>* followed by the classification and releasability information.

b. The second line will contain either:

(a) The text *# filename:<space>* followed by the LDIF file name in format described in paragraph 505 or

(b) The text *# url:<space>* followed by the Universal Resource Locator (URL) of the location on the web site from which the file may be downloaded or

(c) The text # *offline:<space>* followed by the LDIF file name in format described in paragraph 505 of a data file which will be provided to the receiving nation via an offline, out of band mechanism (e.g. hand carried CD-ROM).

c. The third line will contain the text *# type:<space>* followed by the replication type (*initial* for initial, *incremental* for incremental, or *verify* for verify).

d. The fourth line will contain the text *# nation:<space>* followed by the ISO 3166 two letter country code of the sending nation.

e. The (optional) fifth line will contain the text # *area:<space>* followed by a number (0-99) as specified in paragraph 445. If the line is absent, a replication area number of *0* is assumed.

f. The sixth line will contain the text *# sequencenumber:<space>* followed by the current Datafile Sequence Number (DSN)[4].

g. The (optional) seventh line will contain the text *# schemaversion:<space>* followed by a character sting containing a leading alphabetic character followed by one or more digits. These comprise the ACP 133 Edition character followed by the point release number.

h. The (optional) eighth line will contain the text *# format:<space>* followed by the type of data file (*LDIF* for the default LDIF format). If the line is absent, *LDIF* is assumed.

i. The (optional) ninth line will contain the text # compression:<space> followed by the type of compression used (if any). none and ZIP are the only values initially supported. If the line is absent, no compression (none) is assumed.

513. In the case where optional lines are not included, a blank line should not be inserted.

514. There is currently two minor exceptions to the rule that the Message Body text and the Heading in the data file itself are similar. These are as follows:

a. Where a data file is segmented into multiple data files, the Message Heading is only present in the first segment. This is necessary to ensure that the original file can be recreated exactly, thus allowing the checksum to be validated over the whole file.

b. Where a file is to be compressed with ZIP, the filename in each Message Body contains a file type of *.zip* thus corresponding to the file type actually transmitted. The filename within the data file itself (or the first segment thereof), before compression is performed, however reflects the file type of the uncompressed data (i.e. normally *.txt*).

FURTHER EXPLANATORY NOTES

515. The following notes should be used as further guidance when creating Griffin DS replication data:

a. End of line can be marked by *<CR><LF>*, *<LF>* or *<EOF>* characters. Support for other line terminators may be added as and when required.

b. The replication keywords and information should not contain white space characters except where noted.

[4] In the case of a segmented file, the four different fields as described within the message heading text cannot be provided since the checksum cannot be calculated before this information is finalized.

c. The replication information is case insensitive (e.g. GB, gb, or Gb may be used).

d. In all cases, hours are based on a 24 hour clock (e.g. 01= 1AM, 13= 1PM).

e. Format and content may need to be revised based on experience.

EXAMPLE GRIFFIN DS MESSAGE

516. The following is a format and content example of an incremental replication, sending a Secret CAN/US Eyes Only LDIF file with no compression, containing no changes created by Canada to be sent to US on 25 March 2003 at 13:25:59 Zulu time:

File Name:
caus20030325132559Z_0_inc_(SECRET_CANUS_EYES_ONLY).txt

Email and Military Message Subject:
Griffin DS data: caus20030325132559Z_0_inc_(SECRET_CANUS_EYES_ONLY).txt

Message Body Text:
$&#$&#
classification: SECRET CAN/US EYES ONLY
filename: caus20030325132559Z_0_inc_(SECRET_CANUS_EYES_ONLY).txt
type: incremental
nation: CA
area: 0
sequencenumber: 3
schemaversion C0
format: LDIF
compression: none
hash: 78FCE1BE6AB367E1EE30AE951C7701B3DF45E3EE
#&$#&$

LDIF File Attachment Heading:
 # classification: CAN SECRET CAN/US EYES ONLY
 # filename: caus20030325132559Z_0_inc_(SECRET_CANUS_EYES_ONLY).txt
 # type: incremental
 # nation: CA
 # area: 0
 # sequencenumber: 3
 # schemaverson: C0
 # format: LDIF
 # compression: none

CHAPTER 6

IMPLEMENTATION NOTES

BILATERAL AGREEMENTS BETWEEN NATIONS

601. In order to reliably exchange directory information, Nations will need to bilaterally agree how data will be exchanged in each direction. Appendix D lists the top level entries which will be assumed to be present at the receiver by each sending nation. This information may need to differ on a case by case basis. Likewise, Nations must mutually agree to use optional facilities such as subtrees, compression or different data formats. Unilateral adoption by one Nation would not be acceptable.

NATIONAL ARCHITECTURES

602. Although outside the scope of this document's multi-national view of the Griffin DS architecture, each CCEB nation must develop their internal architecture addressing the national interface to the Griffin DS. It is advised that as far as possible, National implementations should be sufficiently flexible to support slight differences in interpretation of this specification and message formats, and should ideally at least to accept without error optional elements on receipt even if there is no intention to send them. Implementations should also be able to handle unforeseen difficulties at the remote end (e.g. need to reset sequence numbers, resend Initial data on request). These mechanisms may require manual intervention.

SPECIAL TOOLS DEVELOPMENT

603. It is envisioned that implementation of the directory solution will require configuration and possibly development of custom tools to satisfy the operational and security requirements of the system. Most instances of these tools would occur within the nations' borders, but would still be essential for a multi-national system operation.

SUPPORT FOR ENHANCED CAPABILITIES

604. Nations must be aware of, and allow for support of various enhancements to the Griffin DS protocol, as described in this document. The most identifiable of these is the support for the ACP 133(C) International Common Content, support for which should be added within a mutually agreed timescale after its ratification by CCEB. Other enhancements may include segmentation and alternative data formats.

605. In many cases, support can be limited to receipt only, since there is no obligation to send data using these new formats and mechanisms. If it is mutually decided by CCEB nations to eventually drop a capability in preference to a better alternative (e.g. LDIF files replaced by DSML files), an agreed interim period (typically a minimum of one year) will be provided to allow time for implementation, trialling and cutover.

BACKWARDS COMPATIBILITY WITH GRIFFIN INTERIM DS

606. Nations must support Griffin Interim DS capability [Ref 6] for interoperability with other nations that do not yet support enhanced capability. This can be achieved by a separate Interim compatible implementation or by optional Interim support within a single implementation. Areas which need to be allowed for backwards compatibility include transport mechanism (SMTP only), Message Control Information (no optional lines) and schema (as defined in Ref 6). Backwards compatibility must be retained until all nations agree that Griffin Interim DS support is no longer required.

ANNEX A TO
CHAPTER 6 TO
ACP 137

SUPPORTED SCHEMA

1. This Annex defines a subset directory schema that is required to support the Griffin DS until it is replaced by a ratified International Common Content within ACP 133(C). The schema consists of attributes which must or may be populated in order to support the Griffin messaging address and contact information requirements. It also indicates the ACP 133 schema type (e.g. structural object class etc) used to incorporate those attributes. It is intended that this Annex will be replaced by ACP 133(C) when a subset of objects and attributes for international interoperability have been ratified.

2. Nations must be capable of defining and receiving information in accordance with this schema. Nations should normally be capable of populating their local directory with not only the "mandatory to populate" attributes, but also other "optional" attributes associated with the schema. However, it is permissible for Nations to filter out unwanted entries or attributes before loading the data. Additionally, LDAP short names for all attributes must be supported.

3. Where attributes are present which require the incorporation of the appropriate auxiliary object class, as defined in ACP 133, it is assumed that the appropriate auxiliary object class will be included.

4. The supported schema may still require some additional work to include certain other CCEB CWAN requirements into organizationalRole and organizationalPerson, which are still under discussion. This document will be modified to reflect these in a future revision.

5. Abbreviations and headings used for schema type include:

   a. **SOC** – Structural Object Class.

   b. **AOC** – Auxiliary Object Class.

   c. **Attr.** – Attribute Definition.

   d. **Mandatory** – Attribute must be populated for this entry.

   e. **Optional** – Attribute need not be populated for this entry.

   f. **Highly Des.** – Attribute is optional but it is highly desirable for this to be populated.

g. **Multivalued** – Attribute may contain more than one value. Griffin Directory Services define that a maximum of four values may be present for any multi-valued attribute, except for the naming attribute of each entry (as indicated in table entry) which is limited to a single value. The Multivalued indicator is preserved for compatibility with ACP 133.

h. **Indexed** – Attribute will be used in search operations and therefore should be indexed to ensure adequate search responsiveness.

i. **Max. Length** – Maximum Length of any attribute value (normally defined as maximum character length). Values specified in parentheses are recommended values suggested for attributes which have no maximum length defined in ACP 133, so that integration testing can be performed.

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max. Length
SOC	X.520: top	Top Superclass	M	
SOC	X.521: country	Country Entry	M	
Populated Attributes				
Attr	X.520: countryName (c)	Name of a country entry in a nation DIT. This is the naming attribute for the entry which should match that indicated in Annex B for each respective Nation.	M	N/N/2

Table A6-1 - Country Entry

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X 501: top	Top Superclass	M	
SOC	X.521: organization	Organization Superclass	M	
SOC	ACP 133: aCPOrganization	ACP 133 Organization Entry	M	
Populated Attributes				
Attr	X.520: organizationName (o)	Name of an organization entry to be held in a Nation's DIT. This is the naming attribute for the entry which should match that indicated in Annex B for each respective Nation.	M	Y/N/64

Table A6-2 - Organization Entry

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X.501: top	Top Superclass	M	
SOC	X.521: locality	Country Entry	M	
SOC	ACP 133: aCPLocality	ACP 133 Locality Entry	M	
Populated Attributes				
Attr	X.520: localityName (l)	Name of a locality entry, used to define the structure between the Organization entry and the Person and/or Role entries within a national DIT. This is the naming attribute for the entry	M	Y/N/128

Table A6-3 - Locality Entry

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X.501: top	Top Superclass	M	
SOC	X 521: organizational Unit	Organisational Unit Superclass	M	
SOC	ACP 133: aCPOrganizationalUnit	APC 133 Organization Unit Entry	M	
AOC	ACP 133: aCPEntryCharacteristics	ACP 133 Entry Characteristics AOC	O	
AOC	ACP 133 aCPPlaUser	ACP 133 PLA User AOC	O	
AOC	ACP 133: aCPOtherContactInformation	ACP 133 Other Contact Information AOC	O	
Populated Attributes				
Attr	X.520: organizationalUnitName (ou)	Name of an organizational unit entry, used to define the structure between the Organization entry and the Person and/or Role entries within a national DIT. This is the naming attribute for the entry.	M	Y/N/64
Attr	X.520: mhsORAddresses	X.400 address(es) to be used for MM.	O	Y[5]/N/(255)
Attr	ACP 133: aCPFunctionalDescription	Function or Task Description associated with the directory entry	O	Y/Y/64
Attr	ACP 133: alternatePLAName	Alternate Plain Language Addresses associated with this organization	O	Y/N/55
Attr	ACP 133: plaNameACP127	Plain Language Address of organization.	O	N/N/55

Table A6-4 - Organizational Unit Entry

[5] Only one mhs-or-addresses value should be present for Military Messaging.

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X.501: top	Top Superclass	M	
SOC	X.521: Person	Person Superclass	M	
SOC	X.521: organizationalPerson	Organizational Person Superclass	M	
SOC	RFC 2798: iNetOrgPerson	iNetOrgPerson Superclass	M	
SOC	ACP 133: aCPOrganizationalPerson	ACP 133 Organization Person Entry	M	
AOC	ACP 133: aCPEntryCharacteristics	ACP 133 Entry Characteristics AOC		
AOC	ACP 133: aCPOtherContactInformation	ACP 133 Other Contact Information AOC		
Populated Attributes				
Attr.	X.520: commonName (cn)	Name of an Organizational person entry within a national DIT. This is the naming attribute for the entry.	M	Y/Y/64
Attr.	X.520: surname (sn)	Last name of person represented by the entry.	M	Y/Y/40
Attr.	X.520: givenName	First name of person represented by the entry.	HD	Y/N/16
Attr.	X.520: telephoneNumber	Telephone number in agreed International format e.g. +1 6139900342.	HD	Y/N/32
Attr.	X.520: organizationalUnitName (ou)	Used to identify the organization the entry is associated with.	HD	Y/Y/64
Attr.	X.520: initials	Initials of person represented by the entry.	O	Y/N/5
Attr.	X.520: facsimileTelephoneNumber	Fax Number in agreed International format – e.g. +1 6139900342.	O	Y/N/32
Attr.	X.520: title	Describes a person's title (e.g. Mr. Mrs. Ms. Etc).	O	Y/N/64

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
Attr.	X.520: description	Additional information deemed necessary to include.	O	Y/N/1024
Attr.	X.520: mhsORAddresses	X.400 address(es) to be used for MM.	O	Y/N(255)
Attr.	ACP 133: aCPFunctionalDescription	Functional or Task Description associated with the directory entry		Y/Y/64
Attr.	ACP 133: militaryIPPhoneNumber	Military IP Phone Number in agreed format.	HD	Y/N/64
Attr.	ACP 133: nationality	Nation sponsoring person entry.	HD	N/N/64
Attr.	ACP 133: rank	Describes a person's military rank.	HD	N/N/32
Attr.	ACP 133: secureFacsimileNumber	Secure Fax Number in agreed format.	HD	Y/N/32
Attr.	ACP 133: secureTelephoneNumber	Secure Phone Number in agreed format.	HD	Y/N/32
Attr.	RFC 1274: mail	Holds multivalued email addresses. Used to evaluate information domain.	O	Y/N/256
Attr:	RFC 2798: displayName	Preferred name of a person to be used when displaying entries. Only single value should be specified in the format : surName<comma><space> givenName/Initials<space>title/rank <space>optionalSourceNationInformation	O	N/N/256

Table A6-5 - Organizational Person Entry

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X.501: top	Top Superclass	M	
SOC	X.501: organizationalRole	Organizational Unit Superclass	M	
SOC	ACP 133: aCPOrganizationalRole	Organization Role Entry	M	
AOC	ACP 133: aCPEntryCharacteristics	ACP 133 Entry Characteristics AOC	O	
AOC	ACP 133: aCPPlaUser	ACP 133 PLA User AOC	O	
AOC	ACP 133: aCPOtherContactInformation	ACP 133 Other Contact Information AOC	O	
Populated Attributes				
Attr.	X.520: commonName (cn)	Name of an organizational role entry within a national DIT. This is the naming attribute for the entry.	M	Y/Y/64
Attr.	X.520: organizationalUnitName	Used to identify the organization the entry is associated with.	HD	Y/Y/64
Attr.	X.520: telephoneNumber	Telephone number in agreed International format – e.g. +1 6139900342.	HD	Y/N/32
Attr.	X.520: description	Additional information deemed necessary to include.	O	Y/N/1024
Attr.	X.520: facsimileTelephoneNumber	Fax Number in agreed International format – e.g. +1 6139900342.	O	Y/N/32
Attr.	X.520: mhsORAddresses	X.400 address(es) to be used for MM.	O	Y/N(255)
Attr.	ACP 133: aCPFunctionalDescription	Function or Task Description associated with the directory entry	O	Y/Y/64
Attr.	ACP133: aCPRoleInformation	Contains information associated with the role occupant, typically in the form Last Name, First Name, Rank	HD	Y/Y/1024

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
		(e.g. Smith John Maj).		
Attr.	ACP 133: alternatePLAName	Alternate Plain Language Addressees associated with this entry.	O	Y/N/55
Attr.	ACP 133: militaryIPPhoneNumber	Military IP Phone Number in agreed format.	HD	Y/N/64
Attr.	ACP 133: nationality	Nation sponsoring role entry.	HD	N/N/64
Attr.	ACP 133: secureFacsimileNumber	Secure Fax Number in agreed format.	HD	Y/N/32
Attr.	ACP 133: secureTelephoneNumber	Secure Phone Number in agreed format.	HD	Y/N/32
Attr.	ACP 133: plaNameACP127	Plain Language Address of Entry	O	N/N/55
Attr.	RFC 1274: mail	Holds multivalued email addresses. Used to evaluate information domain.	O	Y/N/256

Table A6-6 - Organizational Role Entry

Schema Type	Source: Identifier	Description	Mandatory / Optional / Highly Des	Multivalued / Indexed / Max Length
SOC	X.501: top	Top Superclass	M	
SOC	ACP 133: aCPAddressList	Organization Role Entry	M	
AOC	ACP 133: aACPPlaUser	ACP 133 PLA User AOC	O	
Populated Attributes				
Attr.	X.520: commonName (cn)	Name of an addressee list role entry within a national DIT. This is the naming attribute for the entry.	M	Y/Y/64
Attr.	X.520: description	Additional information deemed necessary to include.	O	Y/Y/1024
Attr.	X.520: member	DN of member entry.	O	Y/N(255)
Attr.	X.520: mhs-dl-submit-permissions	List of addresses which may submit using this addressList.	O	Y/N(512)
Attr.	X.520: owner	Specifies the name of some object which has some responsibility for the directory entry that contains this attribute	HD	Y/N/(255)
Attr.	ACP 133: action Addressees	Specifies a list of action addressees of a collective.	O	Y/N/55
Attr.	ACP 133: alternatePLAName	Alternate Plain Language Addresses associated with this AL	O	Y/N/55
Attr.	ACP 133: copyMember	DN of copy member entry.	O	Y/N(255)
Attr.	ACP 133: infoAddressee	Specifies a list of Information Plain Language Addressees of a collective	O	Y/N/55
Attr.	ACP 133: nationality	Nation sponsoring AL entry	HD	N/N/64
Attr.	ACP 133: plaNameACP127	Plain Language Address of AL	O	N/N/55

Table A6-7 - Address List Entry

ANNEX B TO
CHAPTER 6 TO
ACP 137

CCEB NATIONS' HIGH LEVEL DIT STRUCTURES

1. The following table shows the CCEB Nation's High Level DIT Structures.

NATION	High Level DIT
AUS	c=AU, o=GOV, ou=DOD
CAN	c=CA, o=GC, ou=CMIL
NZ	c=NZ, o=NZ Govt, ou=NZDF
UK	c=GB, o=mil
US	c=US, o=U.S. Government

Table 8B1-1 - DIT Structures

REFERENCES

1. Allied Communication Publication 133 (ACP 133) Edition B and Edition C (draft).

2. Federal Information Processing Standard Publication 180-2 (FIPS Pub 180-2). *SHA-1 Algorithm Standard*.

3. RFC 2849, The LDAP Data Interchange Format (LDIF) Technical Specification. *The LDIF format standard.*

4. RFC 3174, US Secure Hash Algorithm 1 (SHA1). Explanation and Example “C” source code for SHA-1 algorithm.

5. Coalition Wide Area Network (CWAN) Working Group (WG) in Section 1 of Tab 2 of Annex C to CWAN-CCEB-DOC-TECH-02-019-VER1.0.0.

6. Griffin Interim Directory Service Technical Architecture Version 1.4 Dated 15 September 2006.

GLOSSARY OF TERMS

ACRONYMS

Acronym	Definition
ACP	Allied Communication Publication
CCEB	Combined Communications Electronics Board
CFBLNet	Combined Federated Battle Lab Network
CR	Carriage Return Character
CWAN	Coalition Wide Area Network
DIT	Directory Information Tree
DN	Distinguished Name
DSML	Directory Services Markup Language
DSN	Datafile Sequence Number
DSWG	Directory Services Working Group
EOL	End of Line
FF	Form Feed Character
FLOT	First Line Of Text
GAL	Global Address List
ISO	International Standards Organization
LDAP	Lightweight Directory Access Protocol
LDIF	LDAP Data Interchange Format
LF	Line Feed Character
LHS	Left Hand Side
MIC	Multinational Interoperability Council
MM	Military Messaging
NATO	North Atlantic Treaty Organization
PKI	Public Key Infrastructure
PLA	Plain Language Address
RHS	Right Hand Side
SMTP	Simple Mail Transport Protocol
SSN	Segment Sequence Number
URL	Universal Resource Locator
UTC	Universal Coordinated Time

DEFINITIONS

Abbreviation	Definition
2-eyes	Network interface between two cooperating CCEB nations.
3-eyes	Network interface between three cooperating CCEB nations.
4-eyes	Network interface between four cooperating CCEB nations.
5-eyes	Network interface between all five CCEB nations.
ACP	Allied Communication Publications – coalition generated technical documents defining mutually-agreed specifications and protocols for interoperability.
CCEB	Combined Communications Electronics Board – coalition composed of Australia, Canada, New Zealand, United Kingdom, and the United States.
CFBLNet	Combined Federated Battle Lab Network -- Coalition classified test network.
DIT	The Directory Information Tree defines the population structure for the directory service.
DSML File	Data file containing directory information in DSML format.
End of Line (EOL)	Character sequence signifying the end of a line of text. This is conventionally Carriage Return (CR) followed by Line Feed (LF) on Windows systems. Other operating systems may employ other EOL conventions including just LF and occasionally Form Feed (FF).
Full Replication	Copy a complete national DIT structure below the agreed replication point.
Griffin	CCEB Coalition operational classified wide area network.
Hash	Typically used for integrity checks. The process of applying a hash function to variable length input to produce a fixed length hash value known as a message digest. The message digest represents concisely the longer message or document from which it was computed. The message digest can be thought of as a "digital fingerprint" of the larger document.
Incremental Replication	Copy only those changes to a national directory that have occurred since the pervious directory exchange.
LDIF file	A common interchange formatted ASCII text file that may contain directory operations and directory data. Typically used to transport directory structure and information between non-connected directory systems.
LHS	The left-hand-side of SMTP email addresses. That part that is mailbox specific and occurs to the left of the @ symbol
Major Domo	An email processor that can be programmed to evaluate email addresses and content (sending and receiving) to readdress, redirect, spawn other mail, and other functions.
Replication	The process of taking a portion (or all) of a directory contents and loading them onto another directory engine.

Abbreviation	Definition
RHS	The right-hand-side of SMTP email addresses. That part that is domain specific and occurs to the right of the @ symbol.
Schema	The Directory schema defines the data structure for a directory application.
SMTP mail	Simple Mail Transfer Protocol is the foundation of most Internet email applications.
X.400	ITU Standard for Message Handling Systems.
Zip File	Data file containing information compression using the ZIP file format.