OPTIONS ANALYSIS OF THE
CANADIAN FORCES PUBLIC KEY INFRASTRUCTURE

ANALYSE D'OPTIONS DE L'INFRASTRUCTURE CLÉ PUBLIQUE DES FORCES
CANADIENNES

A Thesis Submitted
to the Department of Electrical and Computer Engineering
of the Royal Military College of Canada

by

David E DePlanché, CD1
Captain

In Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Computer Engineering
April 2000

© Copyright by D.E. DePlanché 2000

This thesis may be used within the Department of National Defence but copyright for
open publication remains the property of the author.

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-65884-8

Canada

To my wife, Kathleen, and my children Cara and Derek

ACKNOWLEDGMENTS

The author would like to thank Dr. Scott Knight for his help, patience and especially his time in his role as thesis advisor. I'd also like to thank Mr Bob Allen and Major Gareth Gill for their help in the selection of the topic, and Dr. Hafiz Rahman for assisting with network and simulation questions. Finally, thanks to Nick, for setting the pace.

ABSTRACT

DePlanché, David E.. M. Eng (Comp. Eng.). Royal Military College of Canada. April 2000.
Options Analysis of the Canadian Forces Public Key Infrastructure.
Supervisor: Dr. G.S. Knight.

The Department of National Defence (DND) plans to deploy an Entrust® based Public Key Infrastructure (PKI) by the middle of the year 2000. The infrastructure will be deployed on the Defence Wide Area Network (DWAN), a common data communications backbone network interconnecting the heterogeneous Local Area Networks (LANs) and Metropolitan Area Networks (MANs) within the DND.

A PKI uses several cryptographic mechanisms to achieve the following security services: secrecy, authentication, integrity and non-repudiation. These mechanisms include public and secret key cryptography, digital signatures, hash functions, time-stamps, and network security solutions. To effectively examine the impact of the PKI deployment on the DWAN, an operational model of this network was specified. This model was developed using the ComNet network simulation package, and represents a realistic baseline system upon which various PKI configurations are specified.

There are several possible topologies for a PKI, and three have been chosen for detailed analysis. To realize optimal network performance of PKI applications, two variations of these schemes are also presented, combining features that provide the best results during simulations. Each deployment scheme is specified in a different model and PKI message traffic is introduced at increasing levels, representing an increase in the number of applications that use the PKI mechanisms.

The results obtained from the PKI simulations are used to determine the network bandwidth requirements for each specified deployment option. Each scheme is also analyzed to determine the personnel and equipment requirements that would enable a successful deployment of the PKI on the DWAN. The relative life-cycle costs and security risks of each configuration are examined as well, leading to the recommendation of a single PKI deployment scheme.

The use of an operational network model provides a cost effective means to perform several types of configuration and application simulations. Each component of a model can be updated to reflect changes in the underlying technology or in the performance characteristics. This research introduces several PKI deployment schemes, and the associated characteristics and costs, enabling the successful integration of the PKI with the DWAN.

Keywords: Public Key Infrastructure (PKI), timestamp, DWAN, non-repudiation, Entrust, digital signatures, cryptography.

RÉSUMÉ

DePlanché, David E.. M. Eng (Comp. Eng.). Royal Military College of Canada. April 2000.
Analyse d'options de l'Infrastructure Clé Publique des Forces Canadiennes.
Directeur de thèse: Dr. G.S. Knight.

Les Forces canadiennes ont des plans pour déployer une infrastructure à clé publique basée sur Entrust (ICP) pour le milieu de l'an 2000. L'infrastructure sera déployée sur le réseau étendu de la défense (RED), un réseau de base de communications de données de terrain communal interconnectant les réseaux locaux hétérogènes (RLHs) et les réseaux de zone métropolitaine (RZMs) pour le service de la défense nationale (MND).

Une ICP emploie plusieurs mécanismes cryptographiques pour réaliser le secret, l'authentification, l'intégrité et le non-reniement. Ces techniques incluent le cryptography à clé publique et à clé secrète, les signatures digitales, les fonctions d'information parasites, les timbres de temps, et les solutions de sécurité de réseau. Pour examiner pertinemment l'impact du déploiement de ICP sur le RED, un modèle opérationnel de ce réseau a été indiqué. Ce modèle a été développé en utilisant le module de simulation de réseau de ComNet, et représente un système réaliste de ligne de base sur lequel de diverses configurations de ICP sont indiquées.

Il y a plusieurs topologies possibles pour une ICP; trois ont été choisis pour l'analyse détaillée. Pour réaliser une application optimale de ICP, deux variations de ces topologies sont également présentées, combinant les facteurs qui fournissent les meilleurs résultats pendant les simulations. Chaque topologie de déploiement est indiquée dans un modèle différent et le trafic de message de l'ICP est présenté en niveaux croissants, représentant une augmentation du nombre d'applications qui utilisent les mécanismes de ICP.

Les résultats obtenus à partir des simulations de ICP sont employés pour déterminer les conditions sur largeur de bande réseau pour chaque option indiquée. Chaque topologie est également analysée pour déterminer les besoins en personnel et en équipement qui permettraient un déploiement réussi du ICP sur le RED. Les coûts de cycle de vie de parent et les risques de sécurité de chaque configuration sont aussi examinés, menant à la recommandation d'une topologie simple pour le déploiement de ICP.

L'utilisation de modèles opérationnels de réseau fournit un moyen rentable d'exécuter plusieurs types de simulation de configuration et d'application. Chaque composant d'un modèle peut être mis à jour pour refléter des changements dans la technologie fondamentale ou des caractéristiques d'exécution. Cette recherche présente plusieurs topologies de déploiement de l'ICP, et les caractéristiques et les coûts associés, permettant l'intégration réussie du ICP avec le RED.

Mots-clés: L'infrastructure à clé publique (ICP), horodateur, RED, non-reniement, confident, les signatures digitales, cryptography.

TABLE OF CONTENTS

ABSTRACT	v
RÉSUMÉ.....	vi
LIST OF FIGURES.....	xi
LIST OF TABLES.....	xiii
LIST OF ABBREVIATIONS	xiv
1 INTRODUCTION.....	1
1.1 Introduction	1
1.2 Motivation.....	2
1.3 Goals	3
1.3.1 DWAN Traffic and Cost Analysis.....	4
1.3.2 Security Issues.....	4
1.3.3 Time Stamp Issues.....	5
1.3.4 Recommendations for Deployment of PKI and for further Study.....	5
1.4 Thesis Outline.....	6
2 STATE OF THE ART	8
2.1 Introduction.....	8
2.2 Secret Key Cryptography	8
2.3 Public Key Cryptography.....	10
2.4 Combining Secret and Public Key Cryptography	13
2.5 Cryptographic Hash Functions	14
2.6 Digital Signatures.....	14
2.7 User Authentication	16
2.8 Time Stamps.....	17
2.8.1 Network Time Protocol.....	18
2.9 Certificates and Certification Authority	19
2.9.1 CA Cross-Certification.....	21

TABLE OF CONTENTS, CONTINUED

2.10	Certificate Policy and Certificate Practice Statements	22
2.11	DND Certificate Policy and CPS	23
2.12	Security	25
2.12.1	Cryptographic Algorithm Security	26
2.12.2	Hash Functions	27
2.12.3	Time Stamp Security	28
2.12.4	Common Criteria	29
2.12.5	DWAN (Network) Security	33
2.13	DND PKI	36
2.14	Scope of the Present Work	39
3	OPTION SELECTION	40
3.1	Introduction	40
3.2	Criteria for Configuration Selection	40
3.2.1	PKI Traffic	42
3.2.2	Costs	47
3.2.3	Security Risks	48
3.2.4	Complexity and Manageability	48
3.2.5	Criteria Quantification	49
3.3	Configuration Selection	50
3.3.1	Option 1: Centralized PKI Components	53
3.3.2	Option 2: Distributed Level Three PKI Components	55
3.3.3	Option 3: Distributed CAs, DSAs, and TSAs – Level 2	56
3.3.4	Variation 1	58
3.3.5	Variation 2	59
3.4	Summary	59
4	DWAN BASELINE MODEL	61
4.1	Introduction	61
4.2	ComNet III Network Simulation Package	62
4.2.1	ComNet Links	64
4.2.2	ComNet Routers and Switches	65
4.2.3	ComNet Message Flow	65
4.2.4	ComNet Simulation Parameters	67

TABLE OF CONTENTS, CONTINUED

4.3 DWAN Architecture and Protocols	68
4.4 DWAN Baseline Parameters.....	70
4.5 ComNet Model of the DWAN	72
4.5.1 DWAN Model Component Specifications	73
4.5.2 DWAN Baseline Model Simulations	75
4.6 Baseline Simulation Results.....	76
5 PKI DEPLOYMENT MODELS.....	77
5.1 Introduction.....	77
5.2 Option 1 – Centralized Deployment.....	77
5.3 Option 2 Distributed Deployment (Layer 3)	81
5.4 Option 3 Distributed Deployment (Layer 2).....	84
5.5 Variation 1 – Fully Centralized (Minimum Servers).....	86
5.6 Variation 2 – Central CAs and Distributed DSAs & TSAs	88
6 SIMULATION RESULTS FOR PKI MODELS.....	92
6.1 Introduction	92
6.2 Sample Results Discussion	92
6.3 Option 1 - Simulation Results.....	97
6.4 Options 2 and 3 Results.....	101
6.5 Variations 1 and 2.....	104
7 PKI MODEL ANALYSIS.....	107
7.1 Introduction	107
7.2 Simulation Results Analysis.....	107
7.2.1 Option 1 Simulations.....	108
7.2.2 Options 2 and 3 Simulations.....	112
7.2.3 Variation 1	113
7.2.4 Variation 2	115

TABLE OF CONTENTS, CONTINUED

7.3 Cost Implications	115
7.3.1 Option 1 Costs.....	117
7.3.2 Option 2 Costs.....	119
7.3.3 Option 3 Costs.....	121
7.3.4 Variation 1 Deployment Costs	122
7.3.5 Variation 2 Costs	124
7.3.6 Costs Summary.....	125
7.4 PKI Security Analysis	126
7.4.1 Option 1 Security Analysis	129
7.4.2 Options 2 and 3 Security Analysis	130
7.4.3 Variation 1 Security Analysis.....	131
7.4.4 Variation 2 Security Analysis.....	131
7.5 Recommended Deployment Scheme.....	132
8 DISCUSSION AND CONCLUSION.....	134
8.1 Introduction.....	134
8.2 Discussion.....	134
8.2.1 DWAN Baseline Model.....	135
8.2.2 DWAN PKI Models	136
8.2.3 Issues Relating to Entrust PKI Options	137
8.2.4 DND PKI Certificate Use.....	139
8.3 Further Research Areas.....	139
8.4 Conclusion.....	141
LIST OF REFERENCES.....	143
BIBLIOGRAPHY.....	145
APPENDIX A BASELINE RESULTS	148
APPENDIX B SIMULATION RESULTS	155
VITA.....	169

LIST OF FIGURES

<i>Number</i>	<i>Page</i>
Figure 2.1 Secret-Key Cryptography	9
Figure 2.2 Public Key Cryptography	11
Figure 2.3 Common Criteria Evolution	29
Figure 2.4 DND PKI High Level	37
Figure 3.1 Option 1 - Centralized CAs, TSA and DSAs.....	53
Figure 3.2 Option 2 - Distributed Level 3 PKI	55
Figure 3.3 Option 3 - Distributed Level 2 PKI	57
Figure 3.4 Variation 1 - High Level View	58
Figure 3.5 Variation 2 - High Level View	60
Figure 4.1 Sample ComNet Network.....	63
Figure 4.2 DWAN Backbone on ComNet III.....	69
Figure 5.1 Option 1 - Centralized PKI	78
Figure 5.2 Option 2 - Distributed Level 3 Components	83
Figure 5.3 Option 3 - Distributed PKI.....	85
Figure 5.4 Variation 1 - Fully Centralized - Minimum Components	87
Figure 5.5 Variation 2 - Central Segment	89
Figure 5.6 Variation 2 - Distributed Segment	90
Figure 6.1 Sample Link Utilization Graph	93
Figure 6.2 Link Utilization - With Exponential Inter-arrival	94
Figure 6.3 Link Utilization - Without Exponential Inter-arrival.....	95
Figure 6.4 Total Transmitted Frames - Option 1.....	96
Figure 6.5 Option 1 - Borden BAR 2 Link Utilization (West).....	98
Figure 6.6 Option 1 – Tunney’s BAR 1 Link Utilization (West).....	99
Figure 6.7 Option 1 - Borden BAR Link Utilization (East).....	100
Figure 6.8 Option 1 - Tunney’s BAR Link Utilization (East).....	101
Figure 6.9 Option 3 - Borden BAR Link Utilization (West)	102
Figure 6.10 Option 3 - ISAP CPU Utilization (West)	103
Figure 6.11 Option 3 - ISAP CPU Utilization (East)	104
Figure 6.12 Variation 1 and 2 Link Utilization - Halifax	105
Figure 6.13 Variation 1 and 2 Link Utilization – Winnipeg	106
Figure 7.1 Option 1 Results	108
Appendix B	
Figure	
Figure B.1 Option 1 Wainwright, Winnipeg, Shilo, Edmonton - Borden	155
Figure B.2 Option 1 Cold Lake, Yellowknife, Moose Jaw, Suffield - Borden	156
Figure B.3 Option 1 Aldergrove, Esquimalt, Comox, Masset - Borden.....	156
Figure B.4 Option 1 Kingston, Trenton, London, North Bay - Borden	157
Figure B.5 Option 1 Toronto, Petawawa - Borden	157
Figure B.6 Option 1 Wainwright, Winnipeg, Shilo, Edmonton - Tunney's.....	158
Figure B.7 Option 1 Cold Lake, Yellowknife, Suffield, Moose Jaw - Tunney's.....	158
Figure B.8 Option 1 Aldergrove, Masset, Comox, Esquimalt - Tunney's.....	159

LIST OF FIGURES, CONTINUED

Figure B.9	Option 1	Kingston, Trenton, London, North Bay - Tunney's	159
Figure B.10	Option 1	Toronto, Petawawa - Tunney's	160
Figure B.11	Option 1	Greenwood, St John's, Gander, Goose Bay - Tunney's	160
Figure B.12	Option 1	Millcove, Halifax, Shearwater - Tunney's	161
Figure B.13	Option 1	Bagotville, Valcartier, Longue Pointe, Gagetown - Tunney's	161
Figure B.14	Option 3	Shilo, Winnipeg, Edmonton, Edmonton - ISAPs	162
Figure B.15	Option 3	Cold Lake, Suffield, Yellowknife, Moose Jaw - ISAPs	163
Figure B.16	Option 3	Aldergrove, Comox, Masset, Esquimalt - ISAPs	163
Figure B.17	Option 3	Kingston, London, Trenton, North Bay - ISAPs	164
Figure B.18	Option 3	Toronto, Petawawa - ISAPs	164
Figure B.19	Option 3	Greenwood, Gander, St John's, Goose Bay - ISAPs	165
Figure B.20	Option 3	Millcove, Shearwater, Halifax, St JEan - ISAPs	165
Figure B.21	Option 3	Bagotville, Gagetown, Valcartier, Longue Pointe - ISAPs	166
Figure B.22	Variation 1 and 2	Winnipeg, Cold lake, Edmonton, Moose Jaw - Borden	167
Figure B.23	Variation 1 and 2	Esquimalt, Kingston, Trenton, Toronto - Borden	168
Figure B.24	Variation 1 and 2	Greenwood, Shearwater, Halifax, Gagetown - Tunney's	168

LIST OF TABLES

<i>Number</i>	<i>Page</i>
Table 2-1 Assurance Levels	31
Table 3-1 Option Selection Criteria	50
Table 3-2 Final Option Selection Criteria	51
Table 5-1 Option 1 PKI Components	79
Table 7-1 Option 1 Bandwidth Use	109
Table 7-2 Option 1 High Link Utilization	110
Table 7-3 Variation 1 Link Utilization	114
Table 7-4 CF Base PKI Personnel	118
Table 7-5 Option 1 Total Costs	119
Table 7-6 Option 2 Total Costs	120
Table 7-7 Option 3 Total Costs	122
Table 7-8 Variation 1 Total Costs	123
Table 7-9 Variation 2 Total Costs	125
Table 7-10 Deployment Costs Summary	126
Table 7-11 PKI Deployment Summary	132
 Appendix A	
Table	
Table A-1 DWAN Link Capacities and Propagation Delays	148
Table A-2 ISAP CPU Utilization	149
Table A-3 DWAN Baseline Results	150

LIST OF ABBREVIATIONS

ADAM	Automated Data for Aerospace Maintenance
AES	Advanced Encryption Standard
BAR	Border Area Router
bps	bits per second
CA	Certification Authority
CC	Common Criteria
CCF	Canadian Central Facility
CF	Canadian Forces
COTS	Commercial Off-The-Shelf
CPS	Certificate Practice Statement
CPU	Central Processing Unit
CRL	Certificate Revocation List
CSE	Communications Security Establishment
CTCPEC	Canadian Trusted Computer Evaluation Criteria
DDCEI	Director Distributed Computer Engineering and Integration
DES	Data Encryption Standard
DN	Distinguished Name
DND	Department of National Defence
DSA	Directory Server Agent
DSS	Digital Signature Standard
DWAN	Defence Wide Area Network
EAA	Electronic Authorization and Authentication
EAL	Evaluation Assurance Level
EC	Electronic Commerce
ECP	Engineering Change Proposal
FIPS	Federal Information Processing Standard
FRS	Frame Relay Service
FTP	File Transfer Protocol
GOC	Government Of Canada
GOES	Geosynchronous Operational Environmental Satellites
GPS	Global Positioning System (Satellites)
HTTP	Hyper-Text Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IFS	Interim Firewall Service
ISAP	Integrated Services Access Point
ISP	Internet Service Provider
ISSO	Information System Security Officer
ITSEC	Information Technology Security Evaluation Criteria

LIST OF ABBREVIATIONS, CONTINUED

Kbps	Kilo bits per second
LAN	Local Area Network
LRA	Local Registration Authority
MAN	Metropolitan Area Network
Mbps	Mega bits per second
MD-5	Message Digest
MIPS	Millions of Instructions Per Second
MITM	Man-In-The-Middle
MMHS	Military Message Handling System
MOU	Memorandum of Understanding
NCR	National Capital Region
NDHQ	National Defence Headquarters
NIST	National Institute for Standards in Technology
NTP	Network Time Protocol
OID	Object Identifier
OS	Operating System
OSPF	Open Shortest Path First
PKI	Public Key Infrastructure
PMI	Privilege Management Infrastructure
POP	Point Of Presence (Internet)
PSS	Project Security Specialist
QoS	Quality of Service
RA	Registration Authority
RSA	Rivest, Shamir, Adleman
SHA-1	Secure Hash Algorithm
SMTP	Simple Mail Transfer Protocol
TCP/IP	Transmission Control Protocol / Internet Protocol
TCSEC	Trusted Computer System Evaluation Criteria
TSA	Time Stamp Authority
UTC	Universal Coordinated Time
VPN	Virtual Private Network
WAN	Wide Area Network

1 INTRODUCTION

1.1 Introduction

The Canadian Forces plans to deploy an Entrust® based Public Key Infrastructure (PKI) by the middle of the year 2000. The impact that will be felt on the Defence Wide Area Network (DWAN) is, without formal testing, open to speculation.

A PKI forms one part of a network security system and uses several mechanisms to achieve communication and transaction security. Public and secret-key cryptography, digital signatures, hash functions, and time-stamps are all used to protect and verify the integrity of data. Non-repudiation can be achieved using digital signatures and time-stamps, providing proof of origin and time of existence for a document. The PKI must have a comprehensive policy in place so that public keys can be associated with a particular individual or community and specific application classes. The PKI policy specifies how and when an entity may use the PKI services, and describes the lifetimes of all sets of keys.

The heart of a PKI includes a Certificate Authority (CA) and a Directory Server Agent (DSA). The Certificate Authority is the trusted entity within an organization that certifies that a public key belongs to a specified individual. The CA first verifies an entity's identity, generates the necessary key pairs, and then binds the user's identity with a public key into a standard format known as a certificate. These CA validated certificates are stored in a publicly accessible directory or DSA, accessible to all. To enable the extension of trust to an outside agency or

organization, a CA has the ability to exchange certificates with other compatible CAs, forming either a trusted hierarchy or a chain of trust.

A problem associated with the deployment of a PKI on a Wide Area Network (WAN) is finding an optimal network configuration. Each possible solution has a set of advantages and disadvantages in terms of network traffic, human resources, infrastructure costs, and manageability.

The purpose of this research is to explore possible options for the PKI deployment on the DWAN and to analyze them through different metrics. Based on the analysis, recommendations for the DWAN configuration will be made. It is expected that these results may be generalized and applied to other wide area PKI installations.

1.2 Motivation

The Air Force currently uses a database system called Automated Data for Aerospace Maintenance (ADAM) to maintain an electronic log set for each aircraft in the Canadian Forces' (CF) inventory. While everything is stored digitally in this system, current regulations require the printing of each aircraft maintenance document so that handwritten signatures may be applied by qualified technicians. To remove this redundant set of paper-based log sets, the use of digital signatures is required to legally bind a technician to the work performed on an aircraft. At National Defence Headquarters (NDHQ), the group responsible for the implementation of Entrust PKI is Directorate Distributed Computing Engineering and Integration 3-5 (DDCEI), the Department of National Defence (DND) PKI development team. The team at DDCEI indicated that any application, including the ADAM system, that

employs a digital signature scheme must conform to the DND PKI policy and use the Entrust software. Before the DND PKI can be deployed, it should first incorporate a time-stamping subsystem to prevent repudiation of digital signatures.

DDCEI's deployment plan for the PKI has the primary components centrally located at Borden and Tunney's Pasture in Ottawa, but the configuration has not been tested or modeled in any form. It would be interesting and useful to conduct an analysis of the PKI deployment options using a realistic model of the DWAN. The results of such an analysis can be made available to DDCEI before the actual deployment of the PKI, ensuring an effective implementation. An additional benefit of this work is the availability of an accurate DWAN baseline model for research, and the provision of a cost-effective testing environment for future application deployment on the DWAN.

1.3 Goals

There are several possible topologies (physical distributions of the components) for a PKI with time-stamping and of these, three have been chosen for detailed analysis. To realize optimal network performance of the PKI application, two variations of these schemes are also presented, combining features that provide the best results during simulations. The main objective is to determine, for each topology, the characteristics and impacts of the configuration being tested, and these are described in the next sections.

1.3.1 DWAN Traffic and Cost Analysis

Starting with the profile of existing traffic based on current usage data, a baseline model is used to represent current load conditions on the DWAN. The three proposed PKI configurations are then separately modeled and changes in all point-to-point link traffic volumetrics are captured. The new traffic introduced by the PKI includes Directory, CA and Time Stamp Authority (TSA) accesses from remote users all across Canada and from International sites, and the administrative traffic between DSAs, CAs and TSAs. The actual encrypted and/or signed messages are represented by the original baseline traffic, since they would have been sent without the PKI deployment.

The PKI traffic results are used to calculate bandwidth and cost requirements, identify potential bottlenecks, and highlight the advantages and disadvantages of each configuration.

1.3.2 Security Issues

Since the DWAN backbone makes use of leased lines, there are a number of security concerns that have to be examined. Recommendations for the physical security of the PKI equipment, increasing resistance to malicious outside and inside attacks, cryptographic algorithm selection, using secure sources of time, and the protection of PKI management traffic are presented. PKI management traffic includes the transfer of keys on the network, distributing lists of revoked users, and if necessary, time synchronization signals, each of which must be protected using a secure exchange protocol.

1.3.3 Time Stamp Issues

A Time Stamp Authority (TSA) can be configured in several ways. It can be run on the same server machine as the CA or DSA, or as a separate entity somewhere else on the WAN. The system clock can be set to Universal Coordinated Time (UTC) by using a radio receiver, satellite receiver, or a modem. If more than one TSA is present in the PKI configuration, they should have their clocks synchronized with a secure version of the Network Time Protocol (NTP).

The primary TSA should have a redundant source of time, and there should be a backup TSA server on a separate machine, since time-stamps will be a critical factor in determining the legality of a document or transaction. If the primary source of time becomes unavailable for any reason, the clock on the TSA should have a secondary source to provide continuous and accurate time-stamp service to the PKI. A secondary TSA server can be utilized to automatically provide time-stamps if the primary TSA server is offline for any reason. Recommendations for a successful and secure TSA configuration will be presented in detail.

1.3.4 Recommendations for Deployment of PKI and for further Study

All findings will be made available to DDCEI 3-5. The recommendations will provide guidance on the effective deployment of the Entrust PKI on the DWAN. All recommendations will be based on quality of service, cost, security and TSA analysis results.

The possible future use of the DWAN models, and areas available for further research are presented.

1.4 Thesis Outline

The current state-of-the-art in PKI technology is presented in Chapter 2. The basic concepts of secret and public key cryptography are described first, followed by sections on cryptographic hash functions, digital signatures, authentication and time-stamps. A Certification Authority is introduced next, and Certificate Policies and Certificate Practice Statements descriptions follow. The final sections of Chapter 2 provide the reader with an overview of the DND PKI, and introduce PKI security issues.

Chapter 3 describes several criteria used in the selection of the various PKI deployment schemes identified for further analysis. A high level view of the selected configurations is provided.

Chapter 4 introduces the DWAN and the methodology used to model this network. The ComNet network simulation package is described in detail before the DWAN model specifications are presented. The model described in Chapter 4 is used as the baseline model, and all PKI deployment schemes are integrated with this model.

Chapter 5 introduces the PKI deployment models that were described in Chapter 3. A detailed description of each deployment model is provided in Chapter 5. A discussion of the results produced by the PKI model simulations, and a sample of the captured metrics, is found in Chapter 6.

Chapter 7 provides a detailed analysis of the PKI deployments. The results from the simulations are discussed first, followed by sections on the cost implications of each deployment scheme. A security analysis for each deployment comes next, followed finally by a recommendation for a deployment scheme.

The final chapter, Chapter 8, includes several discussion sections, identifies further topics for PKI research, and ends with the conclusion.

2 STATE OF THE ART

2.1 Introduction

A Public Key Infrastructure (PKI) is defined to be a comprehensive system providing public-key encryption and digital signature services (1). It is also described as the set of services that allow the use and management of public-key cryptography and certificates, including keys, certificates, and policies (2). Before the DND PKI can be described, it is necessary to describe each of the mechanisms used within the infrastructure, and how they are used to provide authentication, integrity and confidentiality.

A PKI uses secret and public-key cryptography, digital signatures, hash functions, user authentication, and time-stamping in the provision of its security services. Each of these mechanisms is described in detail in the following sections. The concept of a standard public-key certificate is examined, followed by a section on Certification Authorities (CAs) and the directories that store the certificates. PKI security issues are presented next, followed finally by the PKI policy sections.

2.2 Secret Key Cryptography

Secret-key cryptography is the classical model where both communicating parties hold a common secret key used for both encryption and decryption of data. It can be easily defined

using Figure 2.1 (3), where Alice wants to send a message to Bob, without an eavesdropper (Oscar) being able to determine what the original message was.

Figure 2.1 Secret-Key Cryptography

The key from the key source must be shared by both Alice and Bob over a secure channel before communications can commence. The *encryptor* uses the key to encipher each plaintext character of her message, and then transmits the resulting ciphertext over an insecure channel, where Oscar can observe the message. Since Bob is the only other party with the same key as Alice, only he can decipher the message.

This type of cryptosystem has the disadvantage of requiring previous secure communication of the secret-key between the sending parties. Another drawback is the need for a different secret-key for every pair of users. If Alice is on a Local Area Network with 100 other users, she needs 100 secret-keys to be able to communicate with each of them securely. The system must manage nearly 10,000 keys.

The current secret-key encryption standard is the Data Encryption Standard (DES), which was first adopted for use as a standard for unclassified applications on January 15, 1977.

DES is a block cipher which encrypts and decrypts blocks of 64 bits using a 56 bit key. It was expected to be effective for 10 years, but it has still not been replaced. The most recent review extended the life of DES until 1998. As of January 15 1999, DES is recommended for legacy systems only, with triple-DES being the recommended 'interim' encryption algorithm. Triple-DES simply encrypts data using DES three times, with three different keys, before the ciphertext is transmitted. The 56-bit key used in DES is not considered large enough to be secure and a machine capable of performing an exhaustive key search in about 3.5 hours, costing \$1,000,000 (1993), was proposed by Michael Wiener at CRYPTO '93 (3). A \$250,000 key search computer built by Cryptography Research Inc. broke DES in 56 hours in 1998.

The Advanced Encryption Standard (AES) is the designation of the DES replacement, and the selection has been reduced to five possible algorithms as of August 9 1999. They are MARS, RC6, Rijndael, Twofish and Serpent. The final selection will take place after May 15 2000, the deadline for views to be received on the five finalists. One of the original contenders for the AES was called CAST, and this algorithm is included in the Entrust PKI and has been approved for use within the DND. DES, triple-DES, and RC2 algorithms are also available within Entrust. CAST was not selected as an AES finalist, as it finished eighth overall after performance testing.

2.3 Public Key Cryptography

Public Key cryptography was first used to a large extent in 1977 by Rivest, Shamir and Adleman (4) with the introduction of the RSA Cryptosystem. This system was based on an idea presented by Diffie and Hellman in 1976 (5). The main concept of public key

cryptography is the use of a key pair instead of a secret key. The key pair has a public and private component, where the public key of an intended recipient is used to encrypt information just for the recipient. Only the recipient's private key can decrypt the data. A block diagram showing how a public-key system works is shown in Figure 2.2.

Figure 2.2 Public Key Cryptography

Alice first extracts Bob's public key from a public key directory on an insecure channel. She uses this key to encrypt plaintext and send it to Bob. Only Bob's private key can decrypt this plaintext, and Oscar can eavesdrop on the both the transmission and public key source channels.

The RSA public-key cryptosystem is the algorithm used within the DND PKI, and its security is based on the difficulty of factoring large integers. For an integer to be large enough to be considered secure against factoring algorithms, it needs to have at least 200 decimal digits. RSA works as follows (3):

- Bob generates two large prime numbers, p and q .
- He then computes $n = pq$ and $\phi(n) = (p - 1)(q - 1)$, where $\phi(n)$ represents the number of positive integers less than n and relatively prime to n .
- Bob next chooses a random b ($1 < b < \phi(n)$) such that the greatest common divisor (gcd) of b and $\phi(n)$ is 1.
- Bob next computes $a = b^{-1} \bmod \phi(n)$ using the Euclidean algorithm, which is used to determine if a positive integer has a multiplicative inverse.
- Bob publishes n and b in a directory as his public key.

The two secret prime numbers p and q should each be 100 digits or more to provide security against factoring. When Alice wants to communicate securely with Bob, she encrypts each plaintext character x using Bob's public key and computing $y = x^b \bmod n$. Bob decrypts y using the secret a by computing $x = y^a \bmod n$. Note that the values p , q , and a are kept secret, while the values n and b are public.

Elliptic Curve Cryptography has also been identified as an algorithm for public key cryptography, and the result is a system that has the security of the elliptic curve discrete logarithm problem. It is generally considered that this problem is harder than both the integer factorization problem (RSA) and the discrete logarithm problem modulo a prime. With rapid hardware advances, elliptic curve technology could eliminate the need for single-key ciphers, such as DES, within a PKI (6). Section 2.4 describes how both public and secret key cryptography are used together to overcome the disadvantages of each. Elliptic Curve Cryptography could achieve encryption/decryption speeds comparable to current block cipher performance, thereby eliminating the need to use both types of cryptosystem in a PKI.

There are several other public-key algorithms, but RSA is the only one described in detail in this thesis, since it is the one that is used in the DND's PKI.

2.4 Combining Secret and Public Key Cryptography

Both types of cryptosystem have advantages and disadvantages. Secret-key cryptography requires a secure channel to share a secret key between each pair of users. The number of keys in the system can be equal to the number of users squared, since a different key is needed between each pair of users. A public-key system is slow compared to its secret-key counterpart, due to the complex arithmetic functions that need to be performed. RSA can achieve an encryption rate of 600 Kbps, compared to 1 Gbps for DES. This is based on the most efficient current hardware implementations of each algorithm (3).

A typical PKI makes use of both public and secret-key cryptography to overcome the disadvantages of each and increase system performance. If a message is to be sent to more than one recipient, a separate copy must be encrypted for each person, if public or secret-key cryptography is used alone. The solution to this problem is to generate a one-time secret-key using a secret-key cipher system such as DES, and use it to encrypt the document. This key will be relatively small and it is encrypted using the public key of each intended recipient of the message. Each encrypted secret-key is included in the message header, and the recipients use their private keys to decrypt their copy of the secret key, which is then used to decrypt the message.

2.5 Cryptographic Hash Functions

A cryptographic hash function takes a message of arbitrary length and produces a message digest of a fixed size. This hash function must also be one-way and collision-resistant, meaning that it is virtually impossible to determine the original message from the digest, and that it is computationally infeasible to find two different messages that produce the same message digest (3). The primary purpose of using a hash function is to compress a message into a much smaller, unique message digest that is be used for different PKI services such as digital signatures and time-stamping, which are described in the next sections.

Of all hashing algorithms available to cryptography, SHA-1 (Secure Hash Algorithm) and MD-5 (Message Digest) are the two most common choices. SHA-1 is considered 'stronger' than MD-5 and is therefore the selected algorithm for all hash functions within the DND PKI.

Hashing is used to provide proof of data integrity, since any change in the original text during transmission would produce a different message digest, if the cryptographic hash function is collision-free. Combining this data integrity feature with public-key cryptography produces a common digital signature scheme, described in the next section.

2.6 Digital Signatures

An added benefit of public-key cryptography arises from the fact that the key pair can be used in the reverse order. If Alice encrypts some plaintext using her private key, anyone with access to her public key can verify that the message originated from Alice, since she is the

only entity with access to her private key. This process is described as a digital signature, and it is the method used within a PKI for proof of origin for a transaction.

There are several digital signature algorithms available, and the RSA Signature scheme is described in detail here. Using the same conventions described for the RSA public-key cryptosystem, the RSA signature system works as follows (3):

- Alice encrypts each plaintext character x using her private key and computing $y = x^a \bmod n$.
- Bob decrypts y using the public b and computing $x = y^b \bmod n$.

Since only Alice could have encrypted the plaintext, anyone can verify that the message did indeed originate from her, satisfying the definition of a signature. This signature cannot be forged, but care must be taken with the messages, since copies of the document can't be differentiated from the original.

One of the problems with this scheme is that it is slow, compared to secret-key algorithms, since complicated arithmetic operations are being performed during encryption and decryption, such as modular exponentiation in the case of RSA. Another problem arises from the fact that the signature is at least as long as the original message. The solution to both problems is to use a fast cryptographic hash function to first produce a small message digest, and then apply the digital signature scheme to it.

The signing algorithm used within the DND PKI to encrypt message digests and produce a digital signature is RSA-1024. The Entrust PKI provides a choice between the Digital Signature Standard (DSS) and RSA. With RSA, the public exponent of the cryptography scheme is usually much smaller than the private component. This leads to faster

verification of a signature than signing, but this is a benefit since the document is signed once but may be verified many times.

2.7 User Authentication

Authentication is achieved when one entity has assurance that another entity is who it says it is. The authentication of an entity is separate from any other function performed in a PKI and produces a go/no-go result that allows a user to perform other activities or communications.

Authentication of a user can be based on a single factor or multiple factors from four categories: something you have (smartcard or hardware token), something you know (PIN or password), something you are (thumbprint or retinal scan), and something you do (typing characteristics or handwriting style). Single-factor authentication uses only a single method from the listed categories, while multi-factor authentication uses two or more methods during the process. A bank card being used at a bank machine is an example of two-factor authentication, since the user must have a bank card and a PIN. As more factors are employed in authentication, the ability to impersonate a legitimate entity becomes more difficult.

Authenticating oneself to the local environment doesn't require any service from the PKI. Remote authentication, on the other hand, would use the PKI public key services in what is called a challenge-response protocol. Before the remote user is allowed access to the PKI, he/she can be asked to sign a one-time challenge message with their private signing key. If a signed response is returned, the user authenticated himself/herself, and no sensitive information has been revealed over the network (7).

Within the DND, local user authentication will consist of using a smart card and a PIN. The specific DND PKI Policy section describing remote user authentication procedures has not yet been formalized.

2.8 Time Stamps

A special method of verifying the creation or modification date and time of a digital document is needed. With physical material, the medium itself can be signed and notarized, such as a letter or a research journal; this is not so with documents which remain in their digital form throughout their existence. The time-stamp on a digital document must be applied to the data itself and any tampering, even to one bit, must be detectable. Also, the time-stamp should not be forgeable, i.e. different from the actual date and time (8).

The time-stamp subsystem of the PKI also makes use of a cryptographic hash algorithm. The time-stamp server receives a message digest of some data that a user wants time-stamped. The current date and time is appended to the digest and the server re-hashes the result, and signs it, producing a time-stamp token. The time-stamping system doesn't know the original plaintext of what was time-stamped; it only receives the message digest of some data. The originator and all recipients of a document can verify that a document existed at a given point in time and that it has not been changed in any way. This is verified by the Time Stamp Authority's (TSA) signature on the token. A TSA is the trusted agent within a network that is responsible for the assertion of the proper time. If the originator encrypted and signed the original document and it was subsequently time-stamped, the result is data integrity,

confidentiality and non-repudiation. All of these features are part of the PKI used by the DND.

2.8.1 Network Time Protocol

If multiple Time Stamp servers are to be used in a PKI deployment to support non-repudiation, some means of synchronizing the time between them must be implemented. If it can be shown that numerous time-stamp servers have a different time on the system clock, the question of “what is the correct time” could arise, and the validity of the time-stamp associated with a document, and perhaps the document itself, could be questionable. The standard protocol used to synchronize network clocks is the Network Time Protocol (NTP). This protocol provides the mechanisms necessary to coordinate time between any number of computer systems distributed over a wide area. The time servers operate in a hierarchical master-slave configuration, with the master servers connected to a reliable time source via wire, satellite or radio receiver, or modem (9). The intent of the NTP implementation is to synchronize a computer clock over the Internet by accessing some number of primary servers over Internet paths, and in turn this clock would act as the reference for many local clients. Since the DWAN is not to make use of the Internet reference servers, at least one master must be included on the DWAN, providing a reference to any number of Time Stamp servers. The master system clock is set to Universal Coordinated Time (UTC) using a secure time source and its associated software.

The three main variables at work in the NTP are network delay, measured clock error, and clock offset. Synchronization at the 10 ms level is routinely achieved over WANs with

2000 km distances involved. On a LAN, NTP provides accuracy to the 1 ms level. The NTP time-stamp is a 64-bit unsigned fixed-point number, representing elapsed seconds since 0:00 January 1, 1900. The integer part is in the first 32 bits and the fraction part the last 32, providing a maximum accuracy of about 200 picoseconds (9). A reference time-server with an accuracy granularity of 1 second should meet the needs of the DND.

2.9 Certificates and Certification Authority

A vulnerability of the public-key cryptosystem is the well-known Man-in-the-Middle (MITM) attack. Since the public keys are stored in a location that is accessible to anyone, an adversary might gain control of this site and intercept all message traffic. When a legitimate user requests the public key of an intended recipient, the adversary returns a public key for which he holds the corresponding private key. The originator encrypts his message and sends it, unaware that it will be intercepted and decrypted by the adversary. The original message is encrypted again by the adversary, using the public key of the intended recipient, and neither of the legitimate entities knows about the intrusion. Any traffic in either direction can thus be intercepted, decrypted and read, and re-encrypted properly before being forwarded (10).

The solution to this problem is to make use of a trusted Certification Authority (CA). A certificate in a PKI is the binding of some user information with their public key, signed by the CA. The users must first present themselves to a Registration Authority (RA) and provide proof of their identification before the process of generating certificates can begin. After end-users have validated themselves to the RA, the CA is accessed and a certificate is issued. Information on a certificate includes the user's name (in the form of a Distinguished Name

(DN)), the user's public key, the validity period of the certificate, and the purpose of the public key. A DN uniquely identifies individuals within a domain and is comprised of enough information to ensure uniqueness within an organization, regardless of the number of end-users. The CA's digital signature proves that the certificate has not been tampered with. Certificates are stored on a publicly accessible X.500 directory, called a Directory Server Agent (DSA). A X.500 directory system uses the DNs of the user (or subject) and of the issuer to uniquely identify each certificate. To send an encrypted message, an originator must first acquire the certificates for each intended recipient, and validate these certificates using the CA's public key.

The originator must know and trust one public key for this scheme to work. There must be an out-of-band secure distribution of the CA's public key to the users of the PKI. This one trusted public key is called the trust anchor.

The main components of any PKI are the CA, DSA and the Certificates. The Internet Engineering Taskforce (IETF) PKIX working group is the body that makes the decisions regarding PKI standards. The Certificate standard for a PKI is the X.509 v3 certificate, and it defines data formats and procedures related to the distribution of public keys via certificates that have been digitally signed by a CA.

Each certificate is checked for validity by checking its expiration date and by looking up each certificate in a Certificate Revocation List (CRL), ensuring a certificate hasn't been revoked. A CRL is a data structure that contains a list of revoked certificates and it is signed either by the CA or by another trusted agency and stored at a CRL Distribution Point, usually within a DSA (Z). The public key is then extracted from each certificate and used to encrypt the secret data encryption key that was used to encrypt the document using secret-key

encryption. This key is encrypted for each recipient and, upon receipt of the message, the recipient's private key is used to decrypt the secret key, which is in turn used to decrypt the document. To ensure data integrity and originator authentication, a message digest (hash) is generated by the sender and encrypted with the private key, producing a digital signature. The certificate of the originator can be included with the message to enable validation where a DSA isn't available. The recipient validates the sender's certificate using the issuing CA's public key, and uses the public key from this certificate to decrypt the message digest. Another message digest is calculated locally and compared with the received digest to verify data integrity (11).

2.9.1 CA Cross-Certification

Cross-certification between CAs means that previously unrelated domains can be bound together to allow secure communications between their respective end-users. The mechanism is almost identical to user certification, with the difference being that both the subject and issuer portions of the cross-certificate now contains a CA's identity.

If one of the cross-certified CAs is subordinate to another within the same domain, the certification process is called "intra-domain cross-certification." If the two CAs belong to different domains, "inter-domain cross-certification" is the result. Within the DND, there is a level 1 CA that is cross-certified with level two CAs, and this arrangement is intra-domain. The DND level 1 CA is cross-certified with the Government of Canada (GOC) level 0 CA, and since the DND is part of the GOC, it is still considered intra-domain. If the DND CA is cross-certified with a civilian company for the purpose of conducting e-commerce, this would be considered inter-domain cross-certification.

Cross-certification is called unilateral if CA 1 cross-certifies CA 2 without CA 2 signing the identity and public key of CA 1. If CA 1 and CA 2 cross-certify each other and produce two distinct cross-certificates, this is called mutual cross-certification. Depending on the policy in place for each PKI CA in an inter-domain scenario, full or partial trust may be implemented. For example, one of the CAs may wish to limit the trust extension of cross-certification to only certain individuals, to certain groups, to specific purposes, or to the entire community under the CA (7).

2.10 Certificate Policy and Certificate Practice Statements

A Certificate Policy is a high-level statement of requirements and restrictions associated with the intended use of the certificates issued under the policy (7). A Certificate Practice Statement (CPS) is the detailed description of the internal operating procedures of a CA and PKI that issues certificates (7). The CPS should only be used for internal confidential audits and not for cross-certification purposes. The Certificate Policy is the ideal candidate for use in cross-certification and a pointer to this policy can be put into a certificate.

To distinguish one Certificate Policy from another, each is assigned a globally unique Object Identifier (OID). A cross-certificate may contain a Policy Mapping extension, which is used to permit a policy OID in one domain to be designated equivalent to a policy OID in another domain. An OID is a hierarchical string of integers separated by decimal points. As an example of an OID, 1.2.840.113549.1.1.5 represents SHA-1 with RSA, which could be found in the Signature field of a certificate (7).

The Certificate Policy is established by the Policy Authority for a PKI. Within the DND, the Certificate Policy closely follows the Treasury Board's Certificate Policy, and it is described in detail together with the DND's CPS in the next section.

2.11 DND Certificate Policy and CPS

The implementation of the Department of National Defence (DND) PKI is bound by three documents: Treasury Board PKI Policy (12), Concept of Operations (Con Ops) for the National Defence PKI (13) and the Certificate Practice Statement (CPS) (14).

Treasury Board policy for PKIs states that there are four levels of assurance with respect to the certificates issued; rudimentary, basic, medium and high. Different degrees of effort in validating the identity of an individual named in a certificate and protecting the Certificate Authority (CA) apply to each level. A GOC department's PKI can consist of a single CA or several cross-certified CA's, and the PKI will be used for electronic commerce and communication applications that are PKI-enabled.

Every PKI needs to have in place a Certificate Policy and Certificate Practice Statement before individuals can be issued certificates. A Certificate Policy is a set of rules that binds a certificate to a set of applications or purposes. For example, when a CA issues a certificate to an individual, it means that they have been authenticated according to the rules of the policy and that the individual should be trusted since the CA is trusted. A CPS is the detailed description of how the Certificate Policy is to be implemented and maintained. In other words, the policy states the degree of assurance while the CPS describes how that assurance is attained.

The goal of the Government of Canada's PKI is to reduce costs and increase security in transactions. The PKI is one of the building blocks that will allow secure electronic communications, secure public administration, and secure delivery of services in e-commerce transactions. Paper use should be reduced since digital signatures will be binding when applied to a document or contract. Outside agencies will be able to deal with the government electronically and securely. The CAs will be able to cross-certify with other national and international PKIs, increasing the efficiency and trust in dealing with external sources of information or services.

Before a government department can issue public key certificates, it must first establish a CA with the necessary Certificate Policy or Policies and a CPS. Next, it must sign a Memorandum of Understanding (MOU) and cross-certify their CA with the Canadian Central Facility (CCF), which belongs to the Communications Security Establishment (CSE). The Department of National Defence has done this and may now cross-certify with other government departments, but only through the CCF. If the department wishes to cross-certify with an external CA, it must be approved by the Policy Management Authority of the Government of Canada (GOC) PKI. A repository for certificates and Certificate Revocation Lists must be maintained by the department, and it must be registered with the Registrar of Repositories, the government registrar for repository service providers at the CCF.

Personal information has to be protected in accordance with the *National Archives of Canada Act*, the *Privacy Act*, and the *Access to Information Act*. A department must keep copies of private confidentiality (encryption/decryption) keys for data recovery purposes, protect external subscriber's private keys, and never keep a copy of private signing keys. The main reason for ensuring there is a means of retrieving data is that, while sets of keys expire, the data

may be valid for much longer periods. Another reason for keeping copies of the confidentiality private keys is so that data can be recovered if a user is unable to authenticate themselves because of a lost smartcard or forgotten PIN, and a new set of keys has to be issued. The final reason is a matter of ownership. For example, if an employee develops a product for a company, the company should have the right to the documentation, even after the employee is no longer with the firm. If the private confidentiality keys were not archived, the data would not be retrievable.

Within DND, the Concept of Operations for the department PKI describes the high-level architecture, system objectives and goals, PKI components and services, organizational and physical interfaces, and a general operational outlook. The DND PKI is to be used for the protection of sensitive information, Electronic Commerce (EC) and Electronic Authorization and Authentication (EAA).

2.12 Security

Deployment of a PKI on a WAN involves several aspects of computer security to achieve a level of assurance. The cryptographic algorithms, hash functions, time-stamps, and the associated protocols must offer adequate protection from external and internal threats. A WAN has several points of attack, including the public portion or leased lines of the WAN, and the operating systems of the computer systems. All access points to the network must be from within a trusted sub-net or accessible only through a firewall. This firewall should provide protection for the network from outside connections that could represent a source of

an attack. The set of mechanisms required for network protection will depend on the results of a comprehensive threat/risk assessment process, and will vary for each network.

2.12.1 Cryptographic Algorithm Security

Very few cryptographic algorithms have provable or perfect security. One system that has provable security is the One-time Pad. In this system, the sizes of the plaintext, ciphertext and encryption/decryption key are equal. Every possible key of the required size would need to have the same probability of being used, and for every possible plaintext and corresponding ciphertext, there is a unique key. This cryptosystem is used in military and diplomatic situations on those occasions when absolute secrecy is necessary. Its greatest disadvantage is a result of its design; a key can only be used once, and a new key must be generated and distributed over a secure channel for every message sent. Many contemporary cryptosystems rely on computational security to prevent attacks (3). These are not provably secure, but rely on the computational difficulty of a particular class of arithmetic problem to provide the assurance of security.

The RSA algorithm, used for digital signatures and secure key exchanges within the Entrust PKI, bases its strength on the difficulty of factoring large numbers. Current factoring algorithms can only factor numbers having up to 130 digits within what is considered a reasonable amount of time. The factorization of a 129 digit number required 5000 MIPS-years of computing time donated by over 600 researchers in 1994 (3). It is commonplace today to find RSA systems using at least 1024 bits, or about 308 digits, since the number of bits in a binary representation of an integer is $\log_2 10$ times the number of decimal digits (3).

The security of the DES algorithm has been argued and tested for years. The surest way to attack DES is an exhaustive key search. DES keys are only 56 bits, therefore the corresponding keyspace is 2^{56} . To complete a search of all possible keys, it would take about 24 hours if a machine was built with 10^6 chips such that 10^6 keys could be tested per second, using 1977 technology. The expense of the hardware was the major prohibitive factor for a DES breaking system, but today, a machine can be built that would test 16 encryptions simultaneously, test 50×10^7 keys per second, cost about \$1,000,000, and have an average search time of about 3.5 hours. This is why triple-DES, using a 168 bit key (three 56 bit keys), is recommended for use until the AES is selected.

Within the DND PKI, CAST-128 has been approved only for the protection of designated information within the GOC by the Communications Security Establishment (CSE). The keyspace of CAST-128 is 2^{128} , significantly larger than the keyspace of DES, which is 2^{56} . The successor to this 128 bit block cipher is CAST-256, and it is resistant to all types of attack, including ciphertext only, known plaintext (linear), and chosen plaintext (differential) (15). Using CAST-128, the block size is 128, and the number of available plaintexts for a block is 2^{128} . To perform a linear attack would require approximately 2^{122} known plaintexts, making this type of attack impractical. The number of plaintexts required for a differential attack would be even higher; approximately 2^{140} .

2.12.2 Hash Functions

The two hash functions available on the DND PKI are MD-5 and SHA-1. RSA Laboratories developed MD-5, which produces a 128-bit message digest for any size input.

The chances of two messages having the same message digest is on the order of 2^{64} and the difficulty of finding the correct message given a digest is on the order of 2^{128} . Recalling that a secure hash function must be one-way and collision resistant, MD-5 satisfies the requirement.

SHA-1 was developed by the United States government for use in all federal departments. It produces a 160 bit message digest and is considered to be more collision resistant than MD-5. SHA-1 is the only Federal Information Processing Standards (FIPS) - approved method for secure hashing. Relevant FIPS publications will be discussed in a later section.

2.12.3 Time Stamp Security

The Time Stamp Authority acts as another entity on the PKI with its own set of keys signed and verified by the CA, and communication with the time server should be trusted, since the CA is trusted. The server should be physically located in a secure area, and access to the system clock must be restricted to the TSA. This authority should be the only one with the rights to alter the time on the system clock. Since the time-stamp server(s) of the DND PKI will be using Windows NT as the OS, and the system time is not secure under Windows NT, the use of the local clock as a reference clock is not recommended (16). This means that there should be a secure and trusted external time source, like a GPS receiver. Entrust's Time Stamp package is currently only supported on Windows NT platforms.

2.12.4 Common Criteria

The evolution of computer security standards began in 1983 with the introduction of the Trusted Computer System Evaluation Criteria (TCSEC) or 'Orange Book' in the U.S. By 1993, these criteria became the Federal Criteria in the U.S. and the Canadian Trusted Computer Evaluation Criteria (CTCPEC) was adopted in Canada. The European counterparts in the UK, Germany and France were Memo 3 DTI, ZSEIC, and the B-W-R Book respectively. These three criteria were combined to form the Information Technology Security Evaluation Criteria (ITSEC) in 1991 to promote the recognition of each nation's certifications among the others. Cryptosystems are still evaluated using ITSEC in Europe. The new standard based on the U.S. Federal Criteria (FIPS), ITSEC, and CTCPEC is called the Common Criteria (CC), introduced in 1996. In 1999, the CC officially became ISO 15408 (17). (See Figure 2.3)

Figure 2.3 Common Criteria Evolution

Other important Federal Criteria in the U.S. are FIPS 140-1, 180-1 and 186-1. The National Institute for Standards in Technology (NIST) announced the approval of FIPS 186-1, Digital Signature Standard, on December 15, 1998. FIPS 186-1 allows for the use of either the DSA algorithm or the RSA digital signature algorithm in a cryptosystem. The Secure Hash Standard, SHA-1, is specified in FIPS 180-1, dated the April 17, 1995. FIPS 140-1, Security Requirements for Cryptographic Modules, was approved in 1994 and its replacement, FIPS 140-2, is still in draft form. FIPS 140-1 forms the basis of what are the Federal Criteria, and forms a major part of what has become the Common Criteria (CC). The hope for the CC standard is that it will be acceptable worldwide and every country using PKI certificates will be able to recognize and trust the certificates of other countries.

The CC comprises 638 pages broken into three parts; Introduction, Functional Requirements, and Assurance Requirements. The intent of producing a standardized evaluation criteria, such as the CC, was to allow a certified outside agency to evaluate a vendor's security product against a set of models, rules, and requirements that would define a specific level of security assurance. Once a product receives a certification, trust in its security features has been established, to a certain degree.

There are seven evaluation assurance levels (EAL) detailed in the CC ranging from EAL-1 to EAL-7. EAL-1 and 2 represent a low assurance level, EAL-3 and 4 medium, and EAL-5 and 6 represent high assurance. EAL-7 is for the 'state of the art' and is not generally attainable (17). Table 2-1 shows the comparison between the CC assurance levels and the corresponding levels of ITSEC and TCSEC.

Table 2-1 Assurance Levels

	Assurance Level							
TCSEC	D		C1	C2	B1	B2	B3	A1
ITSEC	E0		E1	E2	E3	E4	E5	E6
CC	EAL-0	EAL-1	EAL-2	EAL-3	EAL-4	EAL-5	EAL-6	EAL-7

The CC is compatible with the (Western) evaluation criteria currently in use, since it represents a combination of most of them. CC certificates are mutually recognized by the U.S., U.K., Germany, Canada, and France. In September of 1999, the Australasian certification scheme 'Gatekeeper' was added to the list (17).

Part 1 of the CC includes a general model of evaluation, CC evaluation results, and a structure for expressing requirements and specifications. Part 2 includes the class, family, and component structures. There are 11 Functional and 7 Assurance classes defined by the CC; as an example, FIA is the Functional class covering Identification and Authentication. Specific families of this class include user identification, user authentication, and specification of secrets. The families are broken down into components and elements such as authentication failures, user attributes, and minimum password sizes. Every element of a component must be successfully evaluated before an EAL can be assigned. For each successively higher EAL, more components need to be verified. For EAL-1, only seven components in seven families of the Assurance classes need to be evaluated. To obtain an EAL-4 rating, 39 components of 23 families must be successfully evaluated. Part 2 of the CC also holds the specific element operation guidelines, the catalogue of functional requirements, and application notes (17).

Part 3 contains the CC Assurance Paradigm, the criteria for the assurance classes, the catalogue of assurance requirements, and finally, the evaluation assurance levels (EALs). In Canada, a vendor can only have its computer security product evaluated by an accredited evaluation and certification authority that has been approved by the Communications Security Establishment (CSE). There are only three accredited organizations in Canada at this time: CGI, DOMUS ITSL, and EWA – Canada Ltd (17).

To have a PKI evaluated and obtain a CC evaluation assurance level, the most important functional classes to be certified would be FCS (cryptographic support), FDP (data protection), FLA (identification and authentication), and FMT (security management). The FTP (trusted path) class would also be important, since this class is concerned with trusted communications paths within a PKI. The Entrust PKI is the first to be evaluated against the CC, and it obtained an evaluation assurance level of EAL-3+ on May 16, 1999. The 'plus' means that all of the criteria for EAL-3 have been met, and the cryptographic module has met the EAL-4 level. This EAL puts the DND (Entrust) PKI (v. 4.0a) at a medium assurance level, and this was the required level for the DND PKI. The evaluation was performed by a team from the U.K. against the ITSEC and converted to the CC. Since the Entrust cryptographic modules were already certified to the FIPS 140-1 standard, most of the effort to obtain this CC rating involved documentation. Entrust has version 5.0 of their PKI ready for release, and it needs to be evaluated against the CC. Any modules that have not been changed between releases will not need to be evaluated, but there must be evidence to this effect (17).

2.12.5 DWAN (Network) Security

The network that hosts a PKI has its own set of security concerns. The DWAN is a relatively secure network since it is currently considered closed to the outside world, having only a limited connection to the Internet. This will not always be the case, since there is a plan to open up the DWAN to full Internet e-mail and more trusted Internet HTTP sites using a firewall system.

2.12.5.1 DWAN Security. The DWAN has three security domains; the unclassified domain with access to the Internet through an EAL-3 level gateway, the designated domain, and the classified domain with a hardware Network Encryption System (NES). The network is TCP/IP router-based and connects each of the base/unit MANs or LANs using a local Integrated Services Access Point (ISAP) router. The Open Shortest Path First (OSPF) routing algorithm is used on the DWAN. The DWAN Extension and Consolidation Project (ECP) (18) has been proposed to bring the DWAN backbone up from PROTECTED A to the SECRET level for the classified domain (19). The ECP also includes an extension to the DWAN to include international and deployed units, and the Internet secure gateways.

Each Internet Point of Presence (POP) will present only a single IP address to the Internet to protect against external attacks on the DND architecture. The Internet secure gateway will make use of Proxies, HTTP, FTP, and SMTP to process service requests before they are forwarded through the POP to the appropriate site server. Replies from the Internet are returned through the gateway to the originator. The plan calls for just two gateways, both at Tunney's Pasture in Ottawa. Anti-virus protection for HTTP, FTP and SMTP will be

incorporated into the gateways. Traps for malicious mobile code will be installed at the gateway and desktop level. Since there are several other applications using bandwidth on the DWAN, a cache server will be installed at each base ISAP to improve performance and help offload DWAN traffic. Users shouldn't expect the same Quality of Service (QoS) that a direct Internet connection provides.

Since the Internet is a strictly unclassified domain, any accessible information on the DWAN must be classified accordingly. Any information posted on the web servers outside of firewall protection would be in the public domain and available for viewing to anyone on the Internet. Processes need to be developed and disseminated to rapidly deal with viruses and Trojan horses before a large-scale infection occurs. The initial trouble site must be blocked from the rest of the DWAN until the problem has been resolved.

Security issues for the ECP are the responsibility of the Project Security Specialist (PSS). The PSS is responsible to the project manager and specific duties include security procedures, security architecture, and security policy. There is also a Project Information System Security Officer (ISSO) who is responsible for the overall security of the project and delivered system, and its Certification and Accreditation. The DWAN ISSO is the liaison between base/unit ISSOs and must ensure that all connections to the DWAN meet established rules (18).

2.12.5.2 Virtual Private Networks (VPNs). A VPN attempts to bring remote access, encryption, security, and the Internet together in a single package. If an employee is able to remotely access company servers using the Internet, costs and time are saved. VPNs are an alternative to dedicated frame-relay and leased-line networks, like the DWAN (20).

The concept of a VPN is that of imitating a private network using a public network. It uses the established public infrastructure, such as the Internet. For this example, an Internet Service Provider (ISP) provides local dial-in connections from various geographic locations, or POPs. By using local POPs, the cost of long distance rates are avoided. Upon connection to a remote ISP, the user's right of access to the company's servers is first verified, then a connection is granted. To ensure secure data transmission, one VPN protocol uses what is called "tunneling". Microsoft's Point-to-Point Tunneling Protocol (PPTP) encrypts the data of a point-to-point packet, including the source and destination data. The encrypted packet is then encapsulated within a Generic Routing Encapsulation (GRE) packet that has the destination address, and the packet is sent (20).

The IPSec protocols and SOCKS are other secure VPN protocols available. PPTP is currently the most widely used VPN protocol, and the user can send TCP/IP packets securely via the Internet. The IETF has proposed IPSec as the suite of protocols used to authenticate users and encrypt data (20).

In the case of a National VPN, like one required to support the DND, an ISP provider would have to be found with POPs at all of the geographic locations hosting a CF base or unit. The provider would also have to be able to handle increased traffic loads and have a scalable product. International traffic would present a problem, since the ISP provider may only operate within the home country. Issues such as multiple vendors, local requirements, and

increased costs would have to be resolved before a VPN could be used to replace the current DWAN deployment (20).

To authenticate a VPN user, a simple user ID and password would not suffice. Two-factor authorization is planned for the CF, and incorporates a smart card and a PIN for every user. All VPN traffic would be encrypted using an existing PKI, such as Entrust.

Within the DND, a VPN will be used for remote access by personnel that require the ability to use e-mail, perform tele-work, or take part in videoconferencing from their home or a remote site. These personnel are identified as being critical to their department's decision-making process, such as CF General Officers that are department heads at NDHQ.

2.13 DND PKI

The DND PKI currently consists of seven (7) NT 4.0 CA servers, hosting Entrust® CA software. The servers are physically located in a secure operational facility at National Defence Headquarters (NDHQ). The combination of the DND PKI hardware, software and personnel operating the PKI at NDHQ is referred to as the DND CA. One of the seven CA servers is designated as the DND PKI Level 1 CA and it is subordinate to the Government of Canada (GOC) PKI Level 0 CA located at the CSE. The Level 1 CA certifies six subordinate CAs, cross-certifies with the Level 0 CA, and will cross-certify with other government departmental CAs in the future. Each of the six Level 2 CAs represents a Defence Wide Area Network (DWAN) domain region and these are the CAs that issue certificates to end-users, publicize public key certificates, and distribute CRLs. There is a Local Registration Authority (LRA) at each base/unit, and their primary responsibility is to provide local support to users by

remote administration. The only functions performed by the LRAs are the secure generation and revocation of certificates on behalf of the Level 2 CAs. This DND PKI configuration is illustrated in Figure 2.4, and it forms the basis for the options analysis of various PKI topologies involving timestamp services.

Figure 2.4 DND PKI High Level

DND has adopted the GOC PKI Medium-Level Assurance Certificate Policy. The DND CPS defines the roles and responsibilities for every entity of the PKI in detail. The DND CPS also expands upon the contents of the GOC PKI Certificate Policy. Medium-level assurance includes information with a designation up to and including Protected B. The DND CPS states that all certificates within the PKI will be X.509 v.3, which is the current

international certificate standard (21). Certificates may be issued or assigned to devices, groups, organizational roles or applications. A Certificate Revocation List (CRL) will be maintained on the X.500 directory at NDHQ, but local copies can be cached to reduce network traffic. All CRLs must be refreshed at least once every 24 hours and the DND CA is not to be held liable for the time between a certificate revocation and the next CRL issuance.

To allow for the future retrieval of information, certificates, CRLs and all keys (except private verification keys) will be archived for at least 30 years. All CA correspondence will be archived for at least 10 years. Every subscriber will need two key-pairs, since one set is for confidentiality (encryption/decryption) and needs to be archived, and the other set is for verification (digital signatures). The verification private key can never be backed up on any system, as this would allow one subscriber to fraudulently assume another's identity. The confidentiality private key that is backed up within a PKI requires two Master Users to be logged on before the key is released. Both of a subscriber's current private keys are to be stored in electronic profiles on a hardware token and both public keys are stored at the CA site.

The algorithms that have been selected for the DND PKI are 1024 bit RSA for all signature and encryption public-key pairs. The CAST algorithm has been selected for the Secure Exchange Protocol (SEP) session keys, which are used to transmit new sets of private keys to users. CAST has also been selected as the algorithm to be used for encryption/decryption of documents.

Key validity periods are 1 year for public encryption and verification keys and certificates, and 6 months for the private signature key. The private decryption key has no expiry so that archived information may be retrieved at any future time. The CA server

operating system (Windows NT) is capable of achieving a TCSEC C2-level rating according to the Communications Security Establishment (CSE) and is trusted to secure DND CA services. Each CA server has three 4 GB tapes for back-ups.

The Con Ops states that the Level 1 and 2 CAs will have X.500 directory systems, a long-term archive, an audit subsystem, and a timestamp subsystem. The primary purpose of the timestamp service is to support non-repudiation services. If a user's certificate is revoked because of a possible key compromise, all prior signatures remain valid and only those documents signed after the time of the possible key compromise should be questioned.

2.14 Scope of the Present Work

The integration of the DWAN and the Entrust PKI requires detailed study to ensure a successful deployment. Each possible configuration scheme would have traffic, cost, and complexity impacts. Network security and the protection of sensitive information have to be attainable and realized.

In order to test various possibilities, an accurate model must first be built that accurately reflects the present-day condition of the DWAN. Once this model is complete, the various PKI arrangements can be added, and the changes documented and measured.

Chapter 3 begins this process by examining the theoretical deployment possibilities for the DND PKI and the selection of the best scenarios for detailed testing and analysis.

3 OPTION SELECTION

3.1 Introduction

There are several possible connection schemes for the DND Entrust PKI using a separate server for each CA, DSA, and TSA. Each scheme has a set of advantages and disadvantages based on criteria consisting of traffic impact on the DWAN, costs in terms of both human and physical resources, security of the PKI servers, and complexity or manageability of the system.

The best possible deployment of the PKI would depend entirely on the elements considered the most critical at the time of deployment. For example, if bandwidth usage is the determining factor, the configuration with the lowest backbone link utilization metrics should be considered. The focus of this chapter is to substantiate the selection of the most feasible PKI deployment configurations so that they can be incorporated into network models. The results of the model simulations provide an indication of how the DWAN will respond to each deployment scheme, and these results help with the determining of the PKI configuration that best meets the selection criteria.

3.2 Criteria for Configuration Selection

There are four criteria used to select suitable PKI deployment schemes in this chapter, and they are DWAN bandwidth utilization, deployment costs, server security, and system

complexity. The first factor to consider for a PKI deployment is the traffic impact on the DWAN. Each CF Base has its own LANs and/or MANs connected to the DWAN backbone through an access point. Each of these heterogeneous LANs and MANs operate at speeds of at least 10 Mbps, while the maximum throughput on any DWAN backbone link is 1.536 Mbps (T1). It is reasonable to conclude that any introduction of PKI message traffic would have a greater impact on the DWAN backbone links than on the higher speed LAN and MAN systems, assuming that the base LANs and MANs are not operating at the point of saturation. The link utilization metric figures available from the DND's Intermapper Web Server indicate that the links between the LANs/MANs and the ISAPs have a very low utilization, averaging about 1.2%.

The second selection criterion for a feasible PKI deployment deals with costs, including the personnel and hardware/software costs. While the bandwidth utilization could also be considered a cost, it is described in its own section.

The third criterion covers the security of the PKI system. The PKI components with the highest risk factor are the CAs, since their compromise would jeopardize the integrity of the entire PKI. The time-stamp server also requires security, but it is not considered as high a risk, since it is the CA that deals with the issuance and revocation of certificates.

The final criterion is concerned with system complexity and manageability. Some deployment schemes can have a minimal impact on the DWAN traffic load, but they would be so complex to manage that they are not a realistic scheme.

The following sections examine each criterion in detail, before the selected deployment schemes are introduced.

3.2.1 PKI Traffic

A primary assumption that must be made before describing the PKI traffic is that any message currently being sent without the introduction of the PKI would also be sent after the PKI is deployed. The difference is that the messages could now have encrypted components. Another assumption is that the PKI has already been initiated, i.e. all users have been issued their PKI credentials and are capable of using the PKI as it was intended, so the traffic load for initial set-up and administration is outside the scope of the models described in the next two chapters. There will still be an amount of PKI administrative traffic associated with the PKI that must be accounted for, and this traffic forms part of the overall PKI traffic captured in the models.

The PKI traffic between the three server types (CA, DSA, TSA) and a user allow for several possible scenarios. The specific PKI messages are now described, and they include key and certificate messages, time-stamp messages, and CRL messages. Once these message types are described, an overall picture of PKI traffic is presented.

3.2.1.1 Keys and Certificates. When a new user is added to the PKI, all necessary keys and certificates are created and stored in the appropriate locations. The encryption key pair and the encryption public key certificate are created by the CA, and the encryption public key certificate is stored in a directory entry (i.e. in the DSA). A copy of the decryption private key and the encryption public key certificate are securely sent to the user's client software (Entrust/Entelligence) and stored with the user's profile. A profile is an encrypted file

containing a user's critical information in a local directory, diskette, or on a cryptographic hardware token such as a smartcard. This information is needed by the client software to allow access to the user's keys. The encryption key pairs are backed up by the CA in a secure database for future data retrieval, and it requires two CA 'Master Users' to retrieve one of these keys.

The directory held on the DSA contains all of the publicly available encryption public-key certificates for every user in the PKI, and no others. The only other type of public certificate is the verification public-key certificate, and these are stored locally with each user's profile.

The verification keys are created by the client software (Entelligence) when a user first creates a profile. Only the user's verification public key is securely sent to the CA, since the private verification key, or signing key, should never be backed up or stored externally from the user's profile. The CA creates a verification public key certificate when it receives the public key and stores a copy securely in an encrypted CA database and returns the certificate to the user client software for storage in the user's profile. When a user signs a file using the private verification key, a copy of the verification public key certificate is included with the signed file, enabling the receiver of a message to verify the source without accessing the public directory.

The primary reason for having two sets of keys is that the private verification key will never be backed up in the PKI. If this critical private key was available to anyone but the rightful owner, a digital signature could be repudiated. The CF will maintain the right to be able to recover encrypted data, but no one will ever be able to masquerade as someone they're

not. All verification public keys are maintained in the secure CA database so that any signed document can be verified at any future time.

3.2.1.2 Time Stamp Authority Traffic. The TSA has its own profile and sets of keys, and simply responds to requests from users to digitally time-stamp some data. A user may choose when to time-stamp some data, or the policy for the DND PKI may be expanded to include a set of rules for TSA usage. When a time-stamp is requested, a message containing several fields is sent to the TSA. The field of primary interest in this message is the message digest of a document, produced by a cryptographic hash function, that the user wishes to time-stamp. A time-stamp token is created using the original document's message digest appended with the TSA's date and time. Next, the token has the same cryptographic hash function applied to it, and a new message digest is produced. The token is then encrypted using the TSA's private verification key, producing a digital signature. The TSA only knows that a particular message digest representing a unique set of data was time-stamped and returned to the user, providing proof of the data's existence at a point in time. The TSA does not know what the actual plaintext message is, and since the hash function is one-way, the original message cannot be recovered from the message digest.

Additional network traffic could come from the use of more than one TSA. If there are multiple time-servers, they may be synchronized using NTP or a similar protocol, and additional network traffic would accrue. If the time-servers do not have any synchronization, then the repudiation problem arises again since different servers could have a different time.

3.2.1.3 Certificate Revocation Lists (CRLs). A CRL is consulted every time a file is encrypted or a signature is verified, implying the storage location of the CRL is important to

traffic analysis. It isn't practical to have only one CRL, especially with large numbers of users, and therefore multiple copies are maintained at various CRL Distribution Points throughout the directory hierarchy. Each certificate has a pointer to its specific CRL Distribution Point so the PKI knows where to look to verify the certificate in question. Within the Entrust PKI, each CRL can contain up to 750 entries, and each CRL can be cached locally in a user's profile to minimize traffic increases due to CRL accesses. The number of CRLs that can be cached is a variable that can be set in a user's initialization file for the client software. A CRL can have a lifetime ranging from 4 to 48 hours, and it is set by the PKI security officer to match the policy requirements. Some PKI policies could require the immediate publication of a new CRL upon a certificate revocation, or permit no caching at all, forcing the retrieval of the appropriate CRL for each certificate verification. The DND PKI policy states that all local copies of a CRL must be less than 24 hours old, and new ones will be created and published in the X.500 directories every 24 hours by the CAs. This CRL updating can be accomplished during off hours to reduce the impact on DWAN traffic, and this assumption is incorporated into the design of the PKI models.

3.2.1.4 End-State PKI Traffic. Referring back to the original assumptions that each user has already been initialized and has a valid set of verification and encryption keys, and that the same messages would be sent with or without a PKI, the Entrust traffic can now be defined. It is important at this point to differentiate between PKI messages and other types of traffic. The messages that are currently being sent on the DWAN includes the messages directly generated by users for applications such as e-mail, accesses to DND information and

web servers, and the transfer of personnel data. PKI messages are generated only when PKI services, such as time-stamps and certificate retrievals, are used.

One source of PKI introduced traffic is a 'time-stamp only' request which occurs when the user wants to time-stamp some data and locally store it. Examples of this situation would include the need to provide evidence for the existence of an electronic journal or log book entry, proof of academic or journalistic propriety for research materials, or legal proof of existence of an idea with the intent of obtaining a patent.

Another type of PKI traffic would involve the encrypting of a document for one or more recipients. A secret-key is used to encrypt the document, and then the public encryption key certificate for each recipient must be accessed and retrieved from the nearest directory. A CRL is consulted for each of the certificates retrieved, and the certificates are used to encrypt the secret key for each recipient. The Entrust PKI header for such a document would contain one encrypted string for each recipient, representing this secret key that only the intended recipients can decrypt using their private decryption keys. A user may encrypt data without selecting a recipient when locally stored files need protection. In this case, no traffic would be introduced onto the DWAN.

A third type of PKI traffic results from the use of digital signatures. Since the message originator's public verification key is included with the signed file, any increases in traffic would result only from accesses to the CRL, initiated by the recipient, for the certificate in question. This process assures the recipient that the originator's certificate has been certified by a CA, and that the certificate hasn't expired or been revoked.

Any combination of time-stamps, signatures and encryption are available within the Entrust PKI, generating several possible traffic scenarios.

Finally, there will be administrative traffic between the PKI components. Local Registration Authorities (LRAs) will be interacting with the CAs on a daily basis to recover users who have forgotten passwords, adding and deleting users, and testing of the PKI status.

All of these traffic types are combined into a single source of data flow for each base or unit before being incorporated into a specific model. The specific details are included in Chapter 5, which describes the PKI deployment models.

3.2.2 Costs

Costs include hardware, software and the personnel required to maintain and administer the PKI system. Bandwidth utilization is also a cost, but since traffic is a separate selection criterion for the deployment topologies, it is not included here. The total system cost implications are presented in Chapter 7. Every PKI configuration presented will have a different cost total, and this value may influence the final deployment selection.

The cost of extra PKI servers and secure time source systems that would be required for certain deployment configurations are also cost factors that must be considered. The DND currently owns 15 systems that are designated for the PKI. Some deployment schemes use many more servers than this, and this added expense will be a factor in determining whether or not a deployment scheme is feasible financially. A direct consequence of using a higher number of PKI servers is an increased personnel cost. For every added CA installation, more personnel are required to administer the system, and these added costs are detailed in Chapter 7 for each selected deployment scheme.

3.2.3 Security Risks

The PKI component with the highest risk factor is a CA. If one of these servers is compromised, the integrity of the entire PKI would be questionable. A malicious user who gains Entrust PKI Master User or Security Officer status by obtaining the necessary password or smartcard/PIN can compromise the system in several ways.

Some of the possible malicious actions that could occur include the creation of fake users, the revocation of valid certificates, the deletion of backed-up data, and a denial of CA services by removing the CA from the network or the removal of critical CA files from the compromised server. Important CA configuration settings can be changed, such as the number of security officers required to perform critical CA functions.

The DSA and TSA servers also have an associated security risk, although the disruption of a single server performing either of these functions could be transparent to the user, since each DSA and TSA has at least one back-up source for the needed service. The biggest threat to these systems would be a denial of service resulting from the particular system's removal from the PKI.

3.2.4 Complexity and Manageability

Complexity of the system applies to the number and type of servers that are needed to realize each design and includes the amount of effort required to administer and synchronize these PKI components. This criterion is tied closely with the cost factor, as the personnel costs are based on the amount of PKI hardware and software that needs to be administered. A

system with only a few centrally located PKI servers is far easier to manage than a deployment scheme that uses distributed components.

3.2.5 Criteria Quantification

To examine each of the criteria described in the previous sections, some form of a rating must be applied to each, so that the criteria can be used in the valid selection of likely PKI deployment schemes.

Each of the four criteria can be assigned a rating of Low, Medium, or High with respect to the other deployment possibilities. Each of these levels must first be defined for the individual criterion before they can be used for configuration selection in the next section.

Starting with PKI traffic, the 'Low' ranking applies to the deployments that would have less than 30% of all PKI messages appearing on the DWAN and the other PKI traffic remaining on the local networks. 'Medium' is used when 30-70% of the PKI messages traverse DWAN backbone links, and 'High' is for those schemes where more than 70% of the PKI traffic uses the DWAN.

For the cost criterion, 'Low' applies to those configurations which use the fewest servers and personnel to administer the PKI. 'Medium' describes a scheme that has some central and some distributed equipment and administrators. 'High' is reserved for those cases where there is a large number of servers and personnel needed for the PKI.

The security risk can be described as 'Low' if the PKI servers are centrally controlled, especially the CAs. 'Medium' is used in those cases where there is a risk to the less-critical PKI

DSAs and TSAs, but the CAs are still secure. 'High' refers to those schemes having the greatest threat to all PKI components, especially the CAs.

The final criterion, complexity and manageability, is rated 'Low' when the system is the easiest to administer and this implies the fewest numbers of PKI servers and personnel are needed. A 'Medium' rating means that more personnel are needed to administer more servers, but the deployment can still be considered manageable. A 'High' ranking for this criterion means that the numbers of personnel and servers needed to deploy and maintain this PKI scheme may be unrealistically high.

3.3 Configuration Selection

The three main components of the PKI are the CA, DSA and Time Stamp servers. Each of these can either be centralized or distributed, which allows for eight possible configurations, depicted in Table 3-1.

Before any options can be eliminated, the relative weight of each criterion must be described. The main servers in Table 3-1 are identified as either being distributed (D) or centralized (C) and the criteria are rated as High (H), Medium (M) or Low (L).

Table 3-1 Option Selection Criteria

Configuration	Traffic	Cost	Risk	Complexity
CA (C), DSA (C), TSA (C)	H	L	L	L
CA (D), DSA (C), TSA (C)	H	M	H	M
CA (C), DSA (D), TSA (C)	M	M	M	M
CA (C), DSA (C), TSA (D)	M	M	M	M
CA (D), DSA (D), TSA (C)	M	H	H	M
CA (D), DSA (C), TSA (D)	M	H	H	H
CA (C), DSA (D), TSA (D)	L	M	M	M
CA (D), DSA (D), TSA (D)	L	H	H	H

In discussions with Director Distributed Computer Engineering and Integration (DDCEI) personnel, it was made apparent that the impact to DWAN backbone utilization for a PKI deployment is the most important criterion to consider. The location of the DSA and TSA servers have the greatest influence on traffic loads, since they will be receiving and sending the majority of PKI messages.

In order to select likely deployment schemes for implementation in PKI models, the number of feasible configurations can be logically reduced. Configurations that have the DSAs and TSAs split between the centralized and distributed DWAN locations can be eliminated due to the number of high and medium criteria rankings resulting from the corresponding deployment scheme.

Another scheme that can be discarded is the one where the CAs are distributed with relation to the DSAs and TSAs. The scheme is counter-intuitive since the CAs should always be either co-located with the other PKI servers, or maintained in a central location. This leaves three possible schemes for further analysis, and these are depicted in Table 3-2.

Table 3-2 Final Option Selection Criteria

Configuration	Traffic	Cost	Risk	Complexity
CA (C), DSA (C), TSA (C)	H	L	L	L
CA (C), DSA (D), TSA (D)	L	M	M	M
CA (D), DSA (D), TSA (D)	L	H	H	H

The third configuration of Table 3-2 doesn't seem like a likely candidate due to the High ratings for three of the criteria. Keeping in mind the importance placed on the traffic loads, this type of configuration merits further study. The first configuration is also examined

in detail, since it has a low rating for the other three criteria, and the impact on the DWAN traffic load could turn out to be less than anticipated. The second configuration is also an ideal candidate for a detailed study, due to its low traffic rating and medium security risk factor.

The selected topologies can now be defined and labeled. Option 1 refers to first configuration described in Table 3-2. All of the PKI servers are centrally located on the DWAN, and the impact on DWAN traffic loads is rated as High.

Since a CA can be part of a hierarchy of CAs, the fully distributed scheme of the last configuration of Table 3-2 can be split into two versions; one with a set of level three components at the bases/units, and the other having the level two components removed from the central locations and appearing at each base/unit instead. Both of these versions of the distributed scheme are examined in detail.

Option 2 represents the last configuration from Table 3-2 where the PKI servers are distributed in a three-level configuration. Option 3 is based on this configuration again, but without the centralized level two servers, i.e. the level two servers are distributed.

Option 1, the fully centralized deployment scheme, can use a variable number of servers to host the PKI. Option 1 refers to the deployment plan currently in place for the DND PKI, which has several CAs and DSAs located at the DWAN central locations. A variation of this scheme is also presented which uses a minimum number of servers to host the PKI, and this deployment configuration is referred to as Variation 1 throughout the remainder of this thesis.

The second configuration from Table 3-2 is also targeted for a detailed examination, but since its traffic load should closely match Option 2 and 3, it is only specified in one model with one PKI traffic load. This configuration is referred to Variation 2 from here on.

The abstract descriptions of the three selected options and two variations are described in the next sections.

3.3.1 Option 1: Centralized PKI Components

Option1 uses a centralized scheme for the PKI servers. This also corresponds to the current plan for the Entrust system deployment by the DND. The complexity, costs and security risk for this option are low, since it uses the fewest servers and extra personnel, and it would be relatively easy to implement. It also produces a large amount of PKI traffic because almost every PKI message would traverse a DWAN backbone link. A high level view of this option is seen in Figure 3.1.

Figure 3.1 Option 1 - Centralized CAs, TSA and DSAs

All of the CAs and DSAs and the TSA are located centrally on the DWAN backbone. While the DWAN is described in detail in the next chapter, the definition of ‘centralized location’ should be clarified at this point. There are two physical locations considered central to the DWAN backbone; CFB Borden and Tunney’s Pasture in Ottawa. Every CF Base has a backbone link to each of these locations.

The Master DSA in Figure 3.1 contains every public-key certificate generated by the PKI. The Shadow DSAs installed with each regional CA hold copies of all of the end-user certificates. Each of these Shadow DSAs have the same entries, i.e. every user’s public certificates are in all of the shadow DSAs.

The level 1 and 2 PKI components presented in Figure 3.1 are all located at either CFB Borden or Tunney’s Pasture. The DSAs for the Pacific, Prairie and Central regions are at CFB Borden, while all of the CAs, the TSA and the remaining DSAs are at Tunney’s Pasture in Ottawa. The reason for modeling this centralized deployment scheme, but with split DSA locations, is that this is the initial planned DND PKI deployment scheme.

The LRAs and end-users are shown at the base level in the figure, with just a representative depiction of a single CF Base shown. The bases are connected to the central locations by DWAN backbone links.

The TSA server is assumed to be connected to one of several possible Universal Coordinated Time (UTC) sources. These possibilities include a satellite (GPS or GOES) receiver, a radio receiver, or an Internet connection.

The LRAs at each of the CF bases/units remotely access their corresponding regional level 2 CAs through the DWAN to administer PKI users. The PKI traffic described in section

3.2.1.4 would all be seen on the DWAN, including time-stamps, public key certificate retrievals, and CRL checking.

3.3.2 Option 2: Distributed Level Three PKI Components

This option takes the network of Option 1 and adds a third layer to the PKI hierarchy. A new CA and Shadow DSA are set-up at each base/unit, and the CAs are cross-certified with their corresponding regional CA's at the centralized locations. Each base would maintain its own PKI components and there would be very little certificate retrieval or CRL traffic seen on the DWAN. Figure 3.2 depicts the high level Option 2 design.

Figure 3.2 Option 2 - Distributed Level 3 PKI

Since there are now PKI servers located at the base level, the figure shows the addition of personnel to administer these components.

All public encryption certificates and CRLs would be locally available in the Shadow DSA. A local TSA would also be used, and it would either be synchronized to the master TSA at Tunney's Pasture, or have its own time source. While this scheme eliminates most of the PKI traffic from the DWAN, the traffic still exists; it is now restricted to the base LANs and MANs.

3.3.3 Option 3: Distributed CAs, DSAs, and TSAs – Level 2

The third option eliminates PKI servers at the central locations, with the exception of the level 1 DND CA and the NDHQ CA (now level two) and its shadow DSA. Each base or unit would again have its own CA, TSA and shadow DSA. The CA's would now be considered level two with respect to the GOC level 0 CA. Cross-certification would be done between these distributed CAs and the DND level 1 CA at Tunney's Pasture. The high level representation of Option 3 is shown in Figure 3.3. This figure only shows a few of the bases/units with their level 2 distributed components.

Figure 3.3 Option 3 - Distributed Level 2 PKI

3.3.4 Variation 1

The Variation 1 PKI deployment scheme is an idea that answers the question: Why are there so many PKI servers needed in Option 1? Recalling that Option 1 is the depiction of the current DND plan to deploy the PKI, reducing the number of PKI servers may have a significant impact on DWAN performance. The main thrust of this variation is to reduce the number of PKI level 2 servers to just two CAs, DSAs, and TSAs. These six servers handle all of the PKI service requests from the end-users at each base. The high-level view of this deployment is presented in Figure 3.4.

Figure 3.4 Variation 1 - High Level View

The difference between Option 1 and Variation 1 is that there is one CA/DSA/TSA to serve the western region of Canada, and another set for the east. These are still level two components, and the level 1 DND PKI CA and Master DSA are still installed. This configuration is modeled using a set of PKI servers at Borden and the other set at Tunney's Pasture in Ottawa.

3.3.5 Variation 2

This variation is derived from configuration F described in Table 3-1. This deployment uses centralized CAs and distributed DSAs and TSAs. The anticipated result from a model based on this design would be a minimal increase in DWAN backbone link utilization. The benefit of increased security for the high-risk CA servers is also realized in this scheme, depicted in Figure 3.5.

Variation 2 uses a CA for the western region of Canada, and one for the eastern. There is also a Master TSA, NDHQ Shadow DSA, and the level 1 PKI servers present at the central location. Each base has its own DSA and TSA, and a sample base is depicted in the figure.

3.4 Summary

The conceptual network configurations that have been presented in this chapter for the likely PKI deployment schemes cover only the abstract topologies. The details for each scheme are provided in Chapter 5, where the deployment models are described and the simulation of each PKI set-up is covered. Before the PKI models can be specified, a baseline

model of the DWAN backbone must be designed which accurately reflects the current status of the DWAN. Chapter 4 presents the design of the DWAN Baseline model in detail, including the incorporation of actual DWAN traffic metrics, link capacities, delays, protocols, and hardware components.

Figure 3.5 Variation 2 - High Level View

4 DWAN BASELINE MODEL

4.1 Introduction

The DWAN provides a common computer and data communications infrastructure for the LANs and MANs within DND. This WAN is designed to provide a high-speed and reliable internetworking communications backbone, supporting the TCP/IP protocol, providing interconnection to existing and planned DND systems. This chapter details the method of modeling this network in such a way as to closely mirror the real network.

To simulate the DWAN with the Entrust PKI deployed, it is necessary to first build and test a model of the existing DWAN backbone. Using the three-step approach for developing a credible model detailed by A. Law and W. Kelton ([22](#)), the first step is to produce a model that seems reasonable to people knowledgeable about the system. For the DWAN, this was accomplished using the Defence Information Network (DIN), which has the most recent technical data about the DWAN. This includes the DWAN ECP ([18](#)), DWAN Architecture Document ([23](#)), and general traffic data. While designing the model of the DWAN backbone, several network specialists were also consulted to verify the model design.

The second step of model development is to quantitatively test the simulation output. The way this is done is by changing the value of a parameter and checking the simulation output for significant changes. The specifics of this step are explained in section 4.5.2.

The third step in developing a realistic model is to compare the model's output data with what would be realistically expected from the existing system. If the two sets of data are

comparable, the model of the existing system is considered “valid”. Since real performance data about the DWAN is available, the model should be a very accurate reflection of the real world. Once the baseline model of the DWAN backbone has been validated, the PKI system components can be incorporated into this model. The simulation results from these modified models would be considered an accurate representation of the performance of the actual system with the same modifications.

To build a working model of the DWAN requires a simulation package and detailed information about all of the network hardware, software, protocols, and architecture. The simulation package is ComNet III v.2.1, a network modeling and simulation software package, produced by the CACI Products Company. Section 4.2 describes in detail the ComNet simulation software and how it is used to model real world networks.

After the ComNet package section, the DWAN architecture and protocols are introduced in section 4.3, followed by the description of the DWAN baseline parameters in section 4.4. Section 4.5 describes the implementation of the DWAN baseline model using ComNet, and the final section, 4.6, presents the results of the baseline simulations.

4.2 ComNet III Network Simulation Package

The various architectures and protocols of a computer network can be implemented in a model using ComNet. This enables the specification of a simulation environment closely matching the real characteristics of the network to be modeled.

This simulation software can emulate the design and operating characteristics of a LAN, MAN, or WAN. Since the DWAN is the network that needs to be modeled, the capabilities of ComNet that allow the specification of a WAN need to be presented.

A WAN is composed of routers, switches, connection lines, and processing nodes. ComNet provides several built-in types of each of these, and also allows the network engineer to specify their own specialized components. The specific parameters for every network feature and component are modifiable by the software package user, and the parameters that apply to the DWAN include router type, link characteristics, and the network protocols.

Figure 4.1 Sample ComNet Network

Figure 4.1 shows a sample of the ComNet components that are used in the models of this thesis. This figure shows two message sources, a message response (destination), routers (nodes), links, and a processing node server. Each of these network components has characteristics that can be set to match an existing or planned network configuration. The

specific capabilities of ComNet are covered next, and the ComNet details applied to the DWAN are presented in section 4.5.

4.2.1 ComNet Links

The links in Figure 4.1 labeled Link 1 and 2 are point-to-point links, and Link 3 is Carrier Sense Multiple Access with Collision Detection (CSMA-CD). A link in ComNet can be designated as one of 20 types, covering everything from LAN links to modem pools to virtual links. The two types that are used in the models of this thesis are point-to-point and CSMA/CD. The parameters of any link that can be set within ComNet include link capacity, propagation delay, error probability, and load.

The link capacity is expressed in bits per second (bps) and ComNet has several built-in values available to choose from. For point-to-point WAN links, the available range is from 56 Kbps up to 1.536 Mbps, which is a T1 line. Higher link capacities can be specified by the network designer. The propagation delay for a link represents the time it takes for a network packet to travel across the network from its source to its destination. This delay is usually specified in milliseconds (ms), since the propagation delay is almost always less than one second and more than one ms. The error probability can be set to whatever the actual value on a real network is, or to a value selected by a network engineer, representing a guaranteed value. For example, the quality of service (QoS) level guaranteed by the commercial lease on the DWAN backbone lines is a factor of 99.5%, representing the minimum probability of successful frame transmission and receipt.

The load of a link can be adjusted within ComNet to simulate background traffic. If the current load on a link is known, it can be set before a simulation is run. This will allow the network designer to add additional capabilities to an existing network and monitor the effect on the current connections. Since a point-to-point link is bi-directional, the load can be set for each direction of traffic flow across a link.

A CSMA/CD link has several parameters as well. The speed of the link is selectable within ComNet to one of several network standards (IEEE 802.3). The subnets of the DWAN use both Ethernet 10-BASE-T and Fast Ethernet 100-BASE-T links, which operate at 10 Mbps and 100 Mbps respectively. The CSMA/CD links also have a propagation delay and error probability associated with them.

4.2.2 ComNet Routers and Switches

The nodes in Figure 4.1 represent either a network router or a switch. ComNet has the parameters for over 300 commercial routers and switches built-in. The only nodes used on the DWAN are Cisco 7500 routers. These routers each have a system bandwidth exceeding 2 gigabits per second. Controllable features include bandwidth utilization, application specific priority assignment, guaranteed quality of service, and multi-point packet delivery.

4.2.3 ComNet Message Flow

The message sources and response icons in Figure 4.1 represent programmable sources and destinations for message traffic within a network. The message sources need to

have arrival times, message sizes, destinations, and packet parameters assigned before a simulation can be run.

The message arrival times can be set to any type of probability distribution including Exponential, Gamma, Poisson, and Weibull. The Exponential distribution is used in each model specified for the DWAN, and it specifies the arrival times of messages to a system that occur at a constant rate. This distribution is widely used to model arrival times of events that follow a Poisson pattern. One parameter is used with the Exponential distribution, and it is called the 'mean', expressed as a value of time. Most messages arrive before or just after the mean value, and this implies that the distribution has a long tail and will occasionally receive a message significantly later than the mean time value. This behavior is very useful in modeling random arrival patterns. Although arrival distributions may be characterized as "Poisson," the Poisson function takes as an argument a time interval and returns an integer which specifies how many arrivals will occur in that interval. In modeling it is more useful to know how much time will elapse before the next arrival. This is why the Exponential distribution is used to actually implement Poisson arrival patterns (22).

The list of possible destinations for each message source can be set to a random value, a weighted-list of destinations, or a single destination. Message sizes can be set to a fixed number of bytes, or to a variable message size based on any of the selectable probability distributions. The packet parameters include the protocol and the packetize time. The packet protocol can be selected from a list of common ones, or defined explicitly. For the DWAN, the protocol used is TCP/IP. The message response attached to the server shown in Figure 4.1 has all of the same adjustable parameters as a message source, with one exception; the destination will always be the message source that sent the original message.

The server also has several parameters that can be set in a ComNet model. These include processing times, delays, buffer sizes, and processor capabilities. For the processing nodes used in the DWAN models, the default parameters are used. They include a single processor, a processing time of 5ms, a pooled buffer size of 100 MBytes, and no priorities set. This means that messages are processed on a first-come first-served basis.

This concludes the discussion of the specific ComNet components of interest to this work. There are several simulation parameters that must also be explained, and they are discussed in the next section.

4.2.4 ComNet Simulation Parameters

ComNet allows several runtime parameters to be set before the actual simulations are run. The most important of these is the packet routing protocol. The DWAN uses the Open Shortest Path First (OSPF) routing algorithm which determines the best route for a packet based on number of hops, link state and reliability. Most users of OSPF employ the Standard Cost Metric (24), which is computed as a function of link bandwidth using the following formula: If bandwidth is less than 100 Mbps, the Cost = $100,000,000 / \text{bandwidth in bps}$. If bandwidth is 100 Mbps or more, the Cost = 1. Thus the cost metric for a T1 line, 1,544,000 bps, is 64. Obviously, the higher the bandwidth, the lower the cost metric. Each packet uses these link costs to arrive at their destinations by the 'cheapest' route.

OSPF abstracts the network routers and links into a directed graph and the arc cost is calculated. The shortest paths are then selected based on lowest costs. Each pair of routers is represented by a pair of arcs, one in each direction, which may be of different weights (24).

One other important runtime parameter that needs to be set is the simulation period. This was set to 3600 seconds (1 hour) with a warm-up of 60 seconds. This means that each simulation represents one-hour of network operation. The warm-up period gets network traffic flowing before measurements begin, to better model the real network, which will always have traffic flowing.

ComNet has the ability to generate a number of reports for each simulation. Many aspects of network performance can be measured and used to generate a report. The two types of report of primary interest in modeling a WAN like the DWAN are link utilization and router performance. The exact report layout will be presented in the section detailing the specifics of the DWAN model, section 4.5.2.

4.3 DWAN Architecture and Protocols

The DWAN physical layout is presented in Figure 4.2. The DWAN backbone consists of 6 Border Area Routers (BARs); three at CFB Borden and three at Tunney's Pasture in Ottawa. Each base within the CF has an Integrated Services Access Point (ISAP) serving as the access points to the DWAN backbone. Both the BARs and ISAPs are Cisco 7500 routers.

There are currently 30 ISAPs connected to the DWAN backbone by leased commercial frame relay service. Each of these ISAPs is connected to one BAR in Borden and one in Ottawa, providing redundant paths for all traffic. Each of the three co-located BARs are inter-connected using 100-BASE-T CSMA/CD Fast Ethernet.

Figure 4.2 DWAN Backbone on ComNet III

The two sets of BARs are connected by T1 service which is installed between BAR 3 in Borden and BAR 3 in Ottawa. Six additional routers are connected directly to the DWAN BARs using a 100-BASE-T connection; three in Borden for their MAN and information (DESC) servers, and three in Ottawa for the NDHQ MAN, International traffic and the main NDHQ (Pearkes) building. The Pearkes building is also connected to CFB Borden by a T1 line represented in Figure 4.2 by link # 65.

Each of the links to the remote bases and units consists of leased bandwidth from commercial sources, with capacities ranging from 56Kbps to T1. The one-way propagation delays range from 10 ms to 120 ms from coast to coast, with the exception of Goose Bay, which has a delay of 600+ ms.

At the base level, each ISAP is connected to the local LANs or a MAN, depending on the size of the base. Remote stations and very small units are connected to the bigger bases by a 56Kbps line.

The DWAN works on the TCP/IP protocol at the network layer and frame relay is used to transport the packets over the leased lines of the backbone. Traffic is routed using the Open Shortest Path First (OSPF) protocol, set to the Cisco 7500 default of lowest cost based on number of hops, link state and reliability.

Frame relay is used over the DWAN, dividing the information into frames or packets. Each frame has an address that the network uses to determine the destination of the frame. The frames travel through a series of switches within the frame relay network and arrive at their destination. Frame Relay Service (FRS) relies on the higher layer transport protocols to perform error recovery and flow control functions. This results in higher throughput and lower delay because much of the network's overhead is eliminated.

4.4 DWAN Baseline Parameters

The ISAP - BAR connections use point-to-point links, and each reflects the actual DWAN link capacities and propagation delays. The Visual Route software package was used to calculate actual DWAN link delays at various times of the day and week. The delays for each of the links were always within a few milliseconds of the calculated average. This average delay for each link is used in the baseline DWAN model links. The ComNet simulation package uses this propagation delay as the mean parameter to an Exponential distribution, meaning that sometimes the delay will be longer or shorter.

Each link also has a utilization factor based on actual DWAN traffic data for the month of September, 1999. Link utilization is the ratio between used bandwidth and total available bandwidth, expressed as a percentage. A necessary assumption is that September 1999 would be a typical month for the DWAN. Traffic data was collected using the network management tool HP Openview, and the maximum traffic volume was captured at 15-minute intervals for every BAR and ISAP port, 24 hours a day. The data consists of link incoming and outgoing traffic volume for each backbone link, expressed in bits per second. Since there is a noticeable drop in traffic volume during the 6 pm - 6 am period, the average hourly traffic volume based on the time-period (local to each ISAP) of 6 am - 6 pm, is used as the baseline link utilization. The traffic increases resulting from the use of the PKI will have the most significant impact on this time period.

Several new applications will soon be deployed on the DWAN, including the CF message system (Military Message Handling System - MMHS), command and control systems, and limited Internet connections. The MMHS traffic is available from the 1997 Message Traffic Volume Report, which provides the 'peak-hour' message transmission and reception figures for every base and unit in the CF. This report is Appendix 3 to Annex C of the DWAN ECP (18), and it specifies to prospective contractors the amount of MMHS traffic that the DWAN must be able to handle. This report states that the average message size is 2 Kbytes without attachments. Since this table provided the number of messages transmitted and received during a peak hour and the message size is known, a data rate can be calculated in bits per second. These figures from the MMHS traffic report are added to the current baseline figures. Some command and control traffic (MCOIN) is already on the DWAN, and is captured in the September 1999 utilization figures. There are also a few hundred World Wide

Web sites available through an Interim Firewall Service (IFS), and this traffic is also reflected in the September 1999 numbers. It is difficult to estimate what the utilization of each link will be when all DWAN applications are up and running. The models can be easily updated to incorporate the actual traffic patterns at a future time. The September 1999 data and the inclusion of the MMHS message traffic provide a reasonable estimate of expected DWAN utilization.

The next section presents the specification of the DWAN baseline network as a ComNet model.

4.5 ComNet Model of the DWAN

Using ComNet III network simulation software, the DWAN backbone can be modeled and simulated. Figure 4.2 illustrates the physical layout of the DWAN, generated from the ComNet DWAN baseline model.

The six BARs are labeled B BAR 1 – 3 and T BAR 1 – 3. As can be seen in the figure, each base or unit on the DWAN has a connection to a BAR located at both CFB Borden and Tunney's Pasture, Ottawa.

The links labeled Link 64 and 67 are the 100-BASE-T interconnections between the three BARs and Link 66 connects Borden BAR 3 and Tunney's BAR 3 with T1 service. Links 1 – 60 connect each of the 30 remote bases to a BAR in Borden and a BAR at Tunney's. Each of the ISAP routers has a message source labeled with two letters, representing the source of all DWAN traffic from the corresponding base or unit.

BAR 3 at Tunney's has the NDHQ MAN, Pearkes building, and the International ISAPs connected to it. Pearkes is the main NDHQ building in Ottawa, and it also has a T1 connection directly to Borden's BAR 1, labeled link 65. The Defence Enterprise Servers (DESC) are situated at CFB Borden, and are classified as information servers. They store the major electronic records of the DND.

4.5.1 DWAN Model Component Specifications

4.5.1.1 DWAN Link Properties. The point-to-point links between each base and the BARs have their own individual characteristics, including propagation delay, capacity, error probability, and load. As an example, Link 1 between Winnipeg and Borden BAR 2 has a propagation delay of 37 ms, an error probability of 0.005 (i.e., the probability that a frame has an error is 0.5 of 1%), a link capacity of 1,536 Mbps (T1), and a utilization or load of 2.66% from Borden to Winnipeg and 2.6% from Winnipeg to Borden. Table A-1 of Appendix A, lists the link capacities and propagation delays for the DWAN backbone connections.

To calculate the link utilization for each connection, the raw incoming and outgoing traffic data collected for each ISAP and BAR on the DWAN for September 1999 is used to calculate the average data rates for each link. The 12 busiest hours from 6:00 am – 6:00 pm (local time) were isolated for each base and the average data rate was calculated for these peak hours. This data was divided by the corresponding link capacity to produce an average peak-hours link utilization figure. In summary, the link utilization numbers represent the expected amount of bandwidth in use for this link on the DWAN during the busiest hours, expressed as a percentage.

4.5.1.2 BARs and ISAPs. The Cisco 7500 routers used as the BARs and ISAPs can process traffic at over 2 Gbps and the device input and output buffers have a 100 MByte capacity. With this amount of throughput, the routers shouldn't reach their operating limits based on DWAN backbone traffic alone. They must also process the traffic moving between the LANs and MANs at each location.

Within ComNet, the Cisco 7500 is one of the selectable devices available for use in network models. All of the router parameters are adjustable, depending on the hardware configuration of the actual routers, and these are set to their defaults.

4.5.1.3 Traffic Sources. For each ISAP on the DWAN backbone, there is a corresponding traffic source, labeled with two letters. For example, the source of all DWAN backbone traffic from CFB Winnipeg is represented by the 'wi' message source within the model depicted in Figure 4.2.

The message sources have several parameters that must be set to emulate an existing network. These include scheduling of messages, session set-up, message sizes, inter-arrival times, destinations, and packet protocols. All of these are set within the model to reflect the actual DWAN network parameters, but all of them, with the exception of the packet protocol, are overridden when link background traffic is being used. Since the baseline traffic metrics have been derived from the available September 1999 network data, this is the mode the DWAN baseline model is set to run in. The packet protocol used, as was previously discussed, is TCP/IP.

4.5.2 DWAN Baseline Model Simulations

Step 3 of the simulation process described in the introduction to this chapter involves the verification of the model's results by comparing them with the real world network expectations. Using the set of September 1999 DWAN traffic data as representative of the actual DWAN traffic load, the baseline model of the DWAN was run five times. This number of simulation runs was selected to verify that the results produced were actually very close to the real DWAN performance for every one of the five simulations.

If a network model parameter isn't changed when using ComNet, every simulation will produce the exact same set of results, using the same random seed each time. To verify the model results, a new seed had to be generated before each of the five simulation runs. This was accomplished by changing the load factor on the CSMA/CD link for the information servers (DESC) located in Borden, shown as link 61 in Figure 4.2. The changing of this link load has no impact on the performance of the backbone links, and it was altered for each of the five DWAN baseline simulations to ensure a new seed was generated.

ComNet has the ability to generate many different kinds of reports for each simulation. The two reports selected for the DWAN baseline model simulation results are the link utilization and the ISAP router processor utilization. The five different sets of results from the simulations of the DWAN baseline were used to produce minimum, maximum, and average link utilization figures. The standard deviation and a 99.9% confidence factor were calculated using all five sets of results.

The confidence factor provides a plus/minus figure that is used to produce the error bars on the graphs of the simulation results for the models simulating PKI deployments on the

DWAN. The confidence factor means that another simulation would produce a result within the confidence range 99.9% of the time. This number was selected to provide assurance that the simulation results were as accurate as possible.

4.6 Baseline Simulation Results

The results of the baseline simulations are presented in Appendix A, and Table A-2 lists each of the ISAPs CPU utilization metrics. The minimum, maximum, and average CPU utilization percentages are listed, and the standard deviation and confidence levels are shown for each ISAP CPU based on five separate simulations.

Table A-3 is a similar list, capturing the link utilization percentages for each link in the DWAN backbone baseline model. These figures are within one percentage point of the figures generated from the actual DWAN traffic data, providing a valid baseline model for testing the PKI configurations

Chapter 5 describes the installation of the PKI components onto the DWAN baseline model, implementing the configurations described in Chapter 3.

5 PKI DEPLOYMENT MODELS

5.1 Introduction

The modeling details of the PKI deployment options presented in Chapter 3 are described in this chapter. The results of the simulations are presented in the next chapter. A description of each model and the individual simulation specifications are provided for each of the basic models that have been implemented. Recall from Chapter 3 the three basic models (Options 1, 2, and 3) which represent the deployment options using a completely centralized system, the addition of a distributed layer, and the removal of the centralized components in a distributed scheme respectively. Two additional models are presented, which are variations of the centralized and distributed configurations. Variation 1 is a completely centralized system with a reduced number of PKI servers. Variation 2 distributes the DSAs and TSAs, while keeping the CAs centralized.

5.2 Option 1 – Centralized Deployment

The first deployment option matches the current CF plan for the implementation of the PKI which uses seven (7) Windows NT 4.0 servers, hosting Entrust® CA software, physically located in a secure operational facility at Tunney's Pasture in Ottawa (NDHQ). The DSAs are also located here, with the exception of the three that serve the western and central

bases. These are located at CFB Borden. This centralized implementation is shown in Figure 5.1. This figure was generated using the ComNet model representation.

Figure 5.1 Option 1 - Centralized PKI

All of the backbone components are identical to the baseline model configuration, and the link, node and message parameters described in the previous chapter are used again in this model. The icon labeled CA1 is the level one CA that is cross-certified with the GOC level zero CA. The other six CAs labeled CA2 – CA7 each serve a geographical area of the DWAN. For each CA there is a DSA located either at Tunney's Pasture or CFB Borden. There is also a Master DSA (Central), matched with CA1. All of the other DSAs shadow the Master (Central) DSA, meaning they contain every certificate generated by the PKI. This redundant and distributed DSA set-up ensures that a valid certificate directory should always be available. There is a single TSA located at Tunney's Pasture to handle all time-stamp requests.

The CA associated with each base on the DWAN is presented in Table 5-1, which shows the CA number, the corresponding DSA, and the bases that use these components.

Table 5-1 Option 1 PKI Components

CA	DSA	Bases
CA 1	Central	None
CA 2	BC	Aldergrove, Masset, Comox, Esquimalt
CA 3	Prairie	Winnipeg, Wainwright, Shilo, Edmonton, Cold Lake, Suffield, Yellowknife, Moose Jaw
CA 4	Ontario	Kingston, London, Trenton, North Bay, Toronto, Petawawa, Borden MAN
CA 5	PQ-NB	St Jean, Bagotville, Gagetwon, Valcartier, Longue Pointe
CA 6	NS-Nfld	Greenwood, Gander, St Johns, Goose Bay, Millcove, Shearwater, Halifax
CA 7	NDHQ	Pearkes, International, NDHQ MAN

The total PKI traffic originating at each base is depicted by the three letter icons at the base of the model and attached to each ISAP. Each of these message sources is configured with a weighted destination list, a message generation frequency, and a message size. Recalling the characteristics of PKI message traffic, presented in section 3.2.1, the possible destinations for PKI messages from each base include a DSA, corresponding CA, and the TSA. The DSA and TSA will see the bulk of the traffic since the CA will only receive requests to recover users or create new ones.

The weighted destination lists for each of the PKI message sources are set with a probability of 50%, 2% and 48% respectively, i.e. there is a 50% chance a generated PKI message will be going to the DSA, 48% to the TSA, and 2% to the CA.

Each of the PKI CAs, DSAs and TSAs has a PKI message response associated with it. When a message is received at any of these PKI servers, they are processed and returned to the originator. The message size selected for the PKI traffic is 2K, and this figure has been chosen based on the fact that time-stamp tokens and requests will be small ($< 1K$), and DSA request sizes will vary, depending upon the number of recipients for a message (typically 1 - 4K).

As was the case with the baseline model, the simulation runtime is set to 3600 seconds, with a 60-second warm-up period to get traffic flow on the lines of the WAN before measurements are taken.

This model was run with five increasing factors of PKI traffic loads, calculated from the MMHS message traffic metrics described in Chapter 4, section 4.4. The relative amount of MMHS traffic generated by each base is an indicator of the relative population size of each CF base. For example, CFB Yellowknife sent a 2 Kbyte message every 1500 seconds during the peak hour for that base. Halifax, on the other hand, sent a message every 15 seconds. Based on

these MMHS numbers for every base, a relative population size can be determined and used for PKI message traffic factors. Using CFB Halifax as an example, the Option 1 PKI model starts with a PKI message session being initiated every 15 seconds, the same interval as the MMHS indicated. This first simulation introduces a new PKI message session every 15 seconds, originating at CFB Halifax, and is called Option 1 - 1X, where 1X represents the multiplier of the message traffic. To increase the PKI traffic, the interval is reduced between PKI message generation times. For CFB Halifax in the Option 1 - 2X model, a PKI message is generated every 7.5 seconds, representing a doubling of PKI traffic for that base. The same process is applied to every base, increasing the PKI message traffic volumes from one times the MMHS traffic volume to five times the volume. This means that there are five separate Option 1 models, labeled 1X to 5X.

The five Option 1 models are used to capture the traffic volume impact to the DWAN with the corresponding increasing levels of PKI traffic. The reports generated are the link utilizations and the ISAP CPU utilizations for each base, and these results are presented in the next chapter.

5.3 Option 2 Distributed Deployment (Layer 3)

The second set of models incorporates a third layer of PKI components physically installed at the base level. The concept for this design is based on a hierarchical ordering of distributed PKI components, where the CAs at each base are at level three in a CA hierarchy. This deployment has the effect of each base acting as its own organization within the PKI, and each of the geographical locations is cross-certified with their level two CA, forming a strict

hierarchical trust model (7). The root CA is the level one CA (CA1) and all end-users hold the public key for this CA as their trust anchor for all certificate verification operations. The intermediate CAs at level two should not issue any certificates to end users, only cross-certification certificates to other CAs, maintaining the concept of a CA hierarchy. The end users are certified by their local CA at each base. To access the public key of an entity certified in another domain, the following process occurs. Entity one's certificate is signed by the local CA. This local CA has a cross-certification certificate signed by a level-two CA. This level 2 CA, in turn, has a cross-certification certificate signed by the root level-one CA. Entity two, who wishes to obtain the public key for entity one, starts with the root CA public key and verifies the cross-certification certificate of the level two CA, extracting a copy of this CA's public key. This key is used to verify the cross-certification certificate of the level three CA where entity one is located. Entity two thus obtains the trusted key for the CA that certified entity one. The public key of this CA is used to verify entity one's certificate, and secure communications can now commence. This certificate verification process is performed by the PKI application and is transparent to the end user.

The physical design for the Option 2 deployment is depicted in Figure 5.2. Note the new PKI components installed at each base on the DWAN. The centralized components are still present, but now have a new role. Each of the directories located at the bases hold copies of all of the PKI certificates to reduce the number of PKI messages that need to traverse the DWAN links. Figure 5.2 shows only the bases and their PKI components. The centralized configuration specified in Option 1, having the level one and two CAs at the central locations of CFB Borden and Tunney's Pasture, remains the same in Option 2 and is not shown in Figure 5.2.

Figure 5.2 Option 2 - Distributed Level 3 Components

At each base, the PKI CA, DSA, and TSA servers are represented by the icons labeled with the three-letter code for each base, followed by '-ca'. For example, CFB Winnipeg's servers are indicated by the 'wpg-ca' icon. The PKI traffic for each base still originates at the message source icons labeled with three letters; 'wpg' for CFB Winnipeg. The majority of PKI messages now stay within each base's physical location. All of these messages pass through the base ISAPs so measurements of the processor utilization for the Cisco routers can capture increasing levels of local PKI traffic.

As described for Option 1, the probability of a PKI message going to the DSA or TSA is 50% and 48% respectively, and 2% for the CA. In this set of models, the DSAs and TSAs are all local, and therefore the traffic stays local. CA accesses can be either local or remote, since the centralized CAs are still present. The link, node, and runtime parameters specified for the baseline model remain the same for Option 2. As for Option 1, five separate models representing incrementing levels of PKI traffic by using traffic multipliers were specified, and simulations executed. The results will be presented in the next chapter. Again, for this set of models, the link utilization and the ISAP CPU utilization were captured for comparison purposes.

5.4 Option 3 Distributed Deployment (Layer 2)

The Option 3 set of models removes the centralized PKI components from Borden and Ottawa, and the primary PKI servers are all located at the bases. This configuration is depicted in Figure 5.3, which shows the removal of the centralized PKI servers that were part of Option 1 and Option 2.

Figure 5.3 Option 3 - Distributed PKI

The configurations found at the bases are the same as specified in Option 2, and are not shown in this figure. Notice the local PKI servers for CFB Borden are represented by the 'bdn-ca' processor connected to Borden BAR 1.

The level one CA (CA1) that is cross-certified with the GOC level-zero CA is still located at Tunney's Pasture in Ottawa. The Master (Central) DSA and TSA are also found here, along with the DSA that services the NDHQ region. All baseline model link, node, and runtime parameters are carried over into Option 3. As was described for Option 1 and 2, there are five models for Option 3, with the increasing PKI traffic levels incorporated by using the traffic multipliers. The simulations were run for each of the Option 3 models, and again, the result data is presented in Chapter 6.

5.5 Variation 1 – Fully Centralized (Minimum Servers)

The centralized deployment scheme described as Option 1 might be improved by reducing the system complexity. Replacing the six regional level two CAs with just two CAs, one for the western region and one for the eastern region, would reduce the system complexity significantly. The same applies to the DSAs and TSAs; install one of each at Borden and Tunney's Pasture in Ottawa. This variation of the centralized deployment is seen in Figure 5.4.

Figure 5.4 Variation 1 - Fully Centralized - Minimum Components

The system is more easily managed, hardware and software costs would be reduced, and security of the PKI components is enhanced since the number of personnel requiring access to the critical administrative functions of the PKI servers is reduced.

CA1, the level one CA, is still present and cross-certified with the GOC level zero CA, and the Central (master) DSA is on the same sub-network. There is an East and West CA, DSA, and TSA to service all of the PKI system requests originating at the bases.

An assumption made for both described deployment variations is that the five times traffic multiplier is a reasonable estimate of the amount of PKI traffic that will eventually be seen on the networks. Instead of five separate models with increasing levels of PKI activity for each variation described, only one model was specified using the five times traffic multiplier.

The same baseline link, node, and runtime parameters are used for this simulation. The difference between this variation and the Option 1-5X model is found in the weighted destination lists for the PKI messages. For each base on the DWAN, all PKI messages will go to either the CFB Borden or Tunney's Pasture PKI servers, depending on the base's geographical location. In this model, Ontario and the western bases use CFB Borden, while NDHQ, Quebec and the eastern bases use Tunney's Pasture. The same measurements were taken as those described for Option 1, and they are presented in Chapter 6.

5.6 Variation 2 – Central CAs and Distributed DSAs & TSAs

To minimize DWAN backbone traffic increases resulting from PKI message traffic, consider a model such that the CAs remain at the central locations and a DSA/TSA pair is located at each base on the DWAN backbone. This set-up puts the PKI servers with the

The destination for virtually all of each base's PKI traffic is local to each ISAP, as was described in the Option 2 and 3 sections previously. The local DSA and TSA are represented by the processors labeled with the three letters for the base, followed by the letters 'dt', representing DSA/TSA. All baseline link, node, and run-time parameters are used in the Variation 2 model, and the traffic multiplier is five. The results are presented in Chapter 6, with all of the other simulations.

6 SIMULATION RESULTS FOR PKI MODELS

6.1 Introduction

The results of running the simulations for each of the 17 PKI models are presented in this chapter. The network parameters that were captured for each simulation are the link utilization for each of the backbone connections, and the CPU utilization of the ISAP routers located at each base. The format of the simulation results for Options 1-3 is Microsoft Excel line graphs. The results for the two architecture variations are shown in Microsoft Excel bar graphs, since they were run at only one level of PKI traffic volume. The simulations provide a basis for comparison of the three options and the two variations of PKI deployments.

Before the graphs are described for each option and variation, a detailed examination of a single sample set of results is presented to provide the necessary insight into the format and significance of the graphs.

6.2 Sample Results Discussion

The sample graph presented in this section was generated using a selection of the actual results from the Option 1 Centralized deployment scheme. In this deployment model, the CAs are all located at Tunney's Pasture in Ottawa with the single time-stamp server, and three DSAs are located at CFB Borden. From the description of the simulation set-up for Option 1, one would expect that the results would show a steady increase in link utilization for

each increasing level of PKI traffic. Figure 6.1 depicts the Option 1 link utilization results for the five incremental levels of PKI traffic for a sample of CF bases.

Figure 6.1 Sample Link Utilization Graph

As the graph depicts, the link utilization for each of the four CF bases generally increases as the level of PKI traffic increases through the five sets of traffic multipliers. The error bars are based on the confidence levels attained in the baseline simulations and are barely discernable. In an ideal network, the slope of lines representing link utilization would be linearly increasing for each level of PKI traffic, as long as the link utilization remained below a level where congestion or blocking would occur. Congestion on a link means that packets will take longer to arrive, and blocking means that resources such as link bandwidth and hardware

buffers cannot handle a particular packet. When blocking occurs, the packet would have to be re-transmitted.

To be able to understand why the lines generally follow a cresting pattern, it is necessary to account for the parameters of the simulations that are designed to match the models closely to the real world. If the exponential inter-arrival times used for each source of PKI traffic at the bases are reset to the fixed value used as the parameter for the exponential functions, the slope of each line should become linear.

Figure 6.2 shows the same results for the same four bases expressed in actual link utilization, instead of the percent link utilization increases depicted in Figure 6.1.

Figure 6.2 - Link Utilization - With Exponential Inter-arrival

As this graph illustrates, the actual link utilization for each CF base is quite low; each is under 6%, even with the maximum amount of PKI traffic introduced onto the DWAN. The trend of each line is the same as in Figure 6.1. Now, the exponential inter-arrival factor is replaced by the fixed inter-arrival value used as the exponential parameter, and the result is shown in Figure 6.3. The slope of the lines has become more linear, but the overall trend for the lines still appears to be one of cresting. CFB Shilo's link capacity and MMHS traffic transmission frequency are so low that the results for this base are very unpredictable. The other three bases are relatively large and send messages with a much higher frequency than Shilo, and these bases are the ones that exhibit the clear cresting trends. This is counter-intuitive and an understanding of this cresting trend is required before the rest of the results can be presented.

Figure 6.3 Link Utilization - Fixed Inter-arrival Time

A method of capturing an overview of the entire system performance is to look at the total number of frames being transmitted during each simulation, and then examining the results for an explanation of the cresting trend seen in Figure 6.3. Working again with the five simulations specified for the Option 1 PKI deployment, the total number of transmitted frames were extracted from the results and they are shown in Figure 6.4, along with the baseline model frame transmission total. As seen in this figure, the cresting trend is apparent, and indicates a limitation of the simulation software (ComNet).

Figure 6.4 - Total Transmitted Frames - Option 1

While the link utilization figures do not approach a level where blocking would occur in a network, the total number of frames being transmitted in the system is increasing, but not

in a linear fashion. The only possible cause of this cresting trend is found in the simulation software limitations. The maximum time a ComNet model can be set to simulate is 3600 seconds, or one hour. This is an ideal selection, since each model emulates peak-hour characteristics of the DWAN. Unfortunately, with the 3600 second time limit, not every base gets to successfully send all of the frames it could, if the software was capable of a longer simulation time. Using the exponential inter-arrival time for the PKI messages of each base to closely match the degree of randomness that will be seen on the real network, the ComNet models end up congesting themselves and running out of time before every PKI message is successfully transmitted and/or received. The impact of this limitation on the results for the simulations is minimal, once it has been identified. The results can be interpreted taking the cresting of the total frames into account. The results for the simulations which have the highest PKI traffic volumes now have to be considered a minimum performance indication for link utilization. The actual utilization would be higher, following the trend of the simulations with the lower PKI traffic volumes, but the simulation results for the deployment schemes still produce an estimate of how the DWAN will actually react to each scheme and traffic volume. The next sections present a sample of the results for each PKI deployment Option and Variation specified. The entire set of results is in Appendix B.

6.3 Option 1 - Simulation Results

The simulation results from the centralized PKI deployment scheme specified by the model described in section 5.2 are presented in this section. Each of the graphs in this section show four bases and their corresponding link utilization. The direction or traffic flow (i.e. to or

from a base) is indicated by the chart title and legend key. Figure 6.5 shows link utilization increases for the CF Bases Cold Lake, Winnipeg, Edmonton, and Moose Jaw on the links to CFB Borden BAR 2. The DSAs for these bases are located at Borden, so a large amount of PKI message traffic will use these links to get to the servers installed there. Most of the rest of the PKI traffic would go to the TSA server located at Tunney's Pasture. The link utilization between these same four bases and Tunney's Pasture BAR 1 is depicted in Figure 6.6.

Figure 6.5 - Option 1 - Borden BAR 2 Link Utilization (West)

Figure 6.6 - Option 1 – Tunney's BAR 1 Link Utilization (West)

Continuing with Option 1, the next set of simulation results represent four CF bases located in the eastern half of Canada. These bases use the CAs, DSAs and TSA located at Tunney's Pasture, so the links to CFB Borden should show a minimal change while the links to Tunney's Pasture should have a significant utilization increase. The first graph, Figure 6.7, shows the link utilization increases on the links from CF bases Halifax, Bagotville, Goose Bay, and Gagetown to the CFB Borden BARs. The change in utilization is within 1.5% of a zero change for each link, indicating the lack of use of these links by PKI messages. The error bars representing the confidence values are now very large compared to the utilization changes.

Figure 6.7 Option 1 - Borden BAR Link Utilization (East)

The links from these same bases to Tunney's Pasture are presented in Figure 6.8, and as anticipated, they show a large increase in link utilization for each of the four bases.

Figure 6.8 - Option 1 - Tunney's BAR Link Utilization (East)

The rest of the results for the Option 1 simulations are available in Appendix B. The analysis of the simulation results for each PKI deployment Option and Variation, and the corresponding impact on the DWAN, is presented in the next chapter.

6.4 Options 2 and 3 Results

The results from the Options 2 and 3 simulations are so closely matched, they are presented together in this section. Option 2 added a third hierarchical layer to the centralized scheme of Option 1 and CAs, DSAs, and TSAs were located at each CF base. Most of the PKI messages go to these servers, and not to the centralized components. In Option 3, the

centralized servers were removed, and each base again has its own set of PKI servers that handle 98% of the PKI messages. The backbone link utilization numbers for these two options are expected to be significantly lower than Option 1, and this can be seen in Figure 6.9, the Option 3 link utilization graph for CF Bases Cold Lake, Edmonton, Winnipeg, and Shilo.

Figure 6.9 - Option 3 - Borden BAR Link Utilization (West)

The eastern bases exhibit a similar result for their link utilization figures. Since there is little change in the DWAN links themselves, the method used to track the changes caused by the PKI message traffic for Option 2 and 3 is the measurement the ISAP CPU utilization for each increasing PKI traffic level. This is an estimate of the effect on the traffic load to each

ISAP. Figure 6.10 is a graph of the increasing ISAP utilization for the same four western CF bases depicted in Figure 6.9. The figures show the impact on the loading of the Cisco 7500 routers is negligible, since they are designed to handle 2 Gbps of traffic. Each estimated traffic volume is much less than the router's capabilities. This graph also indicates that both Winnipeg and Edmonton have a significantly higher volume of message traffic than Cold Lake and Shilo.

Figure 6.10 - Option 3 - ISAP CPU Utilization (West)

Looking at the same set of eastern CF bases as those presented for Option 1, Figure 6.11 shows the ISAP CPU utilization increases for these locations. Halifax is one of the bases on the DWAN that sends a great deal of message traffic, while the other three send a small number of messages.

Figure 6.11 - Option 3 - ISAP CPU Utilization (East)

6.5 Variations 1 and 2

Variation 1 reduced the number of PKI servers used in Option 1 so that there is a CA, DSA, and TSA for the western CF bases and another set for the eastern bases. Variation 2 maintains centralized CAs, while distributing a DSA and TSA to each base. Both Variations were specified in models using the five times traffic volume multiplier only. The best way to present the results of these two simulations is in comparison with the other models. Figure 6.12 shows the Variation 1 and 2 link utilization for CFB Halifax together with the Option 1, 2 and 3 utilization figures, all from the five times traffic multiplier models. The baseline

utilization is also indicated in the left-most bar. The link utilization for Variation 1 (Halifax 4) is almost exactly the same as Option 1 (Halifax 1), while the utilization figure for Variation 2 (Halifax 5) is only slightly more than the baseline level. The results from Options 2 and 3 (Halifax 2 and 3 respectively) indicate the reduced DWAN traffic resulting from distributing the PKI servers to the base level.

Figure 6.12 - Variation 1 and 2 Link Utilization - Halifax

Recall that Halifax is an eastern CF base. This means that the primary link used for PKI traffic is the one to Tunney's Pasture. If the same comparison is performed using a western base, the result should look quite different. The PKI traffic originating from each western base uses both the CFB Borden and Tunney's Pasture links in the Option 1 deployment scheme. (Recall that the DSA for each western base is located at CFB Borden, and

the TSA is located at Tunney's Pasture in Option 1). Figure 6.13 represents the link utilization figures for CFB Winnipeg on the connection between Winnipeg and Borden. The obvious difference between this Figure and Figure 6.12 is the utilization figure for Option 1 (Winnipeg 1) is significantly less than Variation 1 (Winnipeg 4), due to the split set of PKI servers specified in the Option 1 PKI deployment.

Figure 6.13 - Variation 1 and 2 Link Utilization – Winnipeg

The analysis of these results in terms of costs, network performance, complexity, and security is presented in the next chapter. For a complete set of results for all simulations, refer to Appendix B.

7 PKI MODEL ANALYSIS

7.1 Introduction

This chapter provides detailed analysis of the results produced by the simulations performed for each of the PKI models. The impact on the traffic load seen on the DWAN backbone links for each of the options and variations described in Chapter 5 is examined in the first section, 7.2.

After the DWAN performance section, a detailed cost analysis is presented with comparative discussions for the PKI deployment schemes. Next, security issues are covered for each configuration, focusing on the security of the PKI servers.

Finally, a summary of the various options and variations is provided, and each configuration's advantages and disadvantages are listed. The recommended deployment scheme for the DND PKI is identified.

7.2 Simulation Results Analysis

The results of the simulations of the DWAN PKI deployment models that were presented in the last chapter are now analyzed in this section. The two types of output that were captured in the simulations are backbone link utilization and ISAP CPU utilization. The sets of results are examined for each of the PKI deployment configurations and the meaning of the results in terms of bandwidth use and network performance is presented.

7.2.1 Option 1 Simulations

The results produced by the centralized Option 1 simulations are based on backbone link utilization. A sample of one of these sets of results is presented in Figure 7.1, showing the link utilization increases for the connections between four CF Bases and the BARs at Tunney's Pasture in Ottawa.

Figure 7.1 - Option 1 Results

This particular figure shows the link utilization for the links from CF Bases Cold Lake, Edmonton, Winnipeg and Moose Jaw respectively to Tunney's Pasture BAR 1. Each line represents the plot of the percentage change in link utilization for the five increasing levels of PKI traffic. Recall that most of the links between the bases and the BARs currently have a low utilization. One measure of traffic impact is to examine the change in total bandwidth

consumed by the network. Analysis for the Option 1 models can proceed by summing the total DWAN system bandwidth utilization for each increasing level of PKI traffic. These values are then compared with the total bandwidth requirements of the baseline model. This analysis method produces a table of results, Table 7-1, which shows the total bandwidth available on the DWAN, the amount needed for the baseline level of traffic, and the bandwidth used for the baseline augmented by the five levels of PKI traffic introduced by the Option 1 models.

Table 7-1 Option 1 Bandwidth Use

Network Configuration	Bandwidth
DWAN Total	47.568 Mbps
Baseline Model	6.11 Mbps
Option 1 1X Traffic Multiplier	6.516 Mbps
Option 1 2X Traffic Multiplier	6.96 Mbps
Option 1 3X Traffic Multiplier	7.215 Mbps
Option 1 4X Traffic Multiplier	7.503 Mbps
Option 1 5X Traffic Multiplier	7.714 Mbps

The results depicted in Table 7-1 indicate that there is a great deal of bandwidth still available on the DWAN backbone links even after the maximum amount of PKI traffic is introduced. The differences in total utilization are very useful in calculating costs in the next section, but they don't show the whole picture. Each DWAN backbone link must be individually assessed at each level of PKI traffic volume to identify links that have the greatest chance of becoming fully utilized.

The DWAN backbone links that could experience some frame blocking problems as a result of an Option 1 PKI deployment are identified in Table 7-2. Blocking becomes a concern

when a link's utilization reaches Cisco router port maximum speed, and frames may be discarded as the router buffer fills.

This table indicates the links that have a utilization factor greater than 25% at the five times multiplier level for an Option 1 configuration. Recall that these link utilization figures represent average peak-hour usage, and that link utilization increases caused by momentary spikes in traffic patterns would produce a higher percentage of bandwidth utilization. These spikes could use the entire bandwidth available, and cause temporary congestion conditions.

Table 7-2 Option 1 High Link Utilization

Link #	Link Capacity	Link Source	Link Destination	Link Utilization
4	256 Kbps	Tunney's BAR 1	Wainwright	42.6%
12	256 Kbps	Tunney's BAR 1	Suffield	31.63%
12	256 Kbps	Suffield	Tunney's BAR 1	28.74%
6	384 Kbps	Tunney's BAR 1	Shilo	29.86%
14	56 Kbps	Tunney's BAR 1	Yellowknife	25.94%
21	56 Kbps	Masset	Borden BAR 1	44.14%
22	56 Kbps	Masset	Tunney's BAR 3	41.07%
30	384 Kbps	Trenton	Tunney's BAR 2	25.11%
30	384 Kbps	Tunney's BAR 2	Trenton	33.7%
56	768 Kbps	Tunney's BAR 1	Gagetown	26.03%
66	1.536 Mbps	Borden BAR 3	Tunney's BAR 3	50.73%
66	1.536 Mbps	Tunney's BAR 3	Borden BAR 3	28.97%
65	1.536 Mbps	Pearkes BAR	Borden BAR 1	31.71%
65	1.536 Mbps	Borden BAR 1	Pearkes BAR	43.94%

The backbone links with the highest utilization for Option 1 – 5X are the links to CF Bases with low available bandwidth. Increasing PKI introduced traffic has a greater impact on these lines. The last four links are presented in the table for the purpose of identifying throughput bottlenecks on the DWAN. The high link utilization indicted for these connections are primarily a result of the current baseline message traffic and these links would not normally be used for PKI messages. If the primary link from a base to a BAR is not available because of the OSPF routing algorithm properties, or if a link or component fails, then link 66 would start carrying some PKI traffic. For example, if the link between CFB Masset and Borden BAR 1 went down, the PKI messages normally using this link would now have to be routed via Tunney's Pasture BAR 3, and from there to the Borden BARs using link # 66, which would show an increase in its utilization. This fault-testing scenario can easily be simulated using the ComNet models. Such testing is outside of the scope of this thesis, and remains as future work.

None of the backbone links between the bases and a BAR exceed 50% utilization with the current level of baseline or Option 1–5X PKI traffic. The baseline traffic load will increase in the future as more applications make use of the DWAN, and the links identified in Table 7-2 may need upgrading to prevent blocking.

The cost calculation for the bandwidth utilization estimated by the set of Option 1 models is not trivial. A standard unit cost for bandwidth on the many DWAN links does not exist. A cost modeling figure provided by DWAN life cycle managers will be used here. The cost of 1.536 Mbps of throughput (T1) is \$930/month. Applying this cost to Table 7-1, the Option 1 configuration will increase the cost of bandwidth utilization by \$245/month for the model using the one times traffic multiplier. The cost will increase to \$970/month for the

Option 1 model using the five times traffic multiplier, which rolls up to an annual cost increase of \$11,640. The total is rounded up to \$12,000 in the remainder of this chapter. This sum does not seem like a great deal of money, but it must be included as a planned expense by some government department. This is a conservative estimate, as it does not account for the cost effect of increasing the bandwidth on multiple links. Increasing the available bandwidth several times to accommodate the increasing link demand may be more expensive than a single, larger increase. For example, upgrading all of the DWAN backbone links that have a bandwidth less than a 1.536 Mbps to a T1, and upgrading all of the T1 lines to a T2 (6.312 Mbps) or T3 (44.736 Mbps) (24), may be more cost effective than increasing each link's capacity by one channel (56 Kbps) at a time. The bandwidth cost also does not account for upgrades to the communication infrastructure to accommodate the required bandwidth.

7.2.2 Options 2 and 3 Simulations

The bandwidth needs for the Options 2 and 3 configurations are significantly less than Option 1. The model for Option 2 using the five times traffic multiplier uses a total of 6.325 Mbps, and the Option 3 model at the same PKI level of traffic needs only 6.157 Mbps. The change in bandwidth costs for these two sets of models can be considered inconsequential in the total cost picture.

The local traffic load increase seen on the LANs and MANs at each CF Base will be the same as the local traffic load increase indicated by the Option 1 configuration. A user will normally be situated on these networks. The PKI traffic generated and received by a user's workstation will be the same for every equivalent level of PKI traffic multiplier used for each

option or variation deployment scheme. An assumption is made that the local networks have a higher capacity and can handle the increased traffic loads resulting from a PKI deployment scheme. The link utilization to the ISAPs from each of the CF Base's LANs averages less than 1.5% each. These links carry all of the traffic between the base LANs and the base ISAP, and include the traffic that is sent from one base LAN to another. This means that the base ISAPs are not operating near full capacity.

The reports generated by the simulations for the Option 2 and 3 sets of models include the ISAP CPU utilization changes resulting from the PKI traffic load increase. The increase in CPU utilization resulting from PKI message traffic is less than .2 of 1% for each base. This means that there is no communications infrastructure related increase in costs for deploying Option 2 and 3. The biggest cost factors for these two deployment options will come from equipment and personnel needs, and these are detailed in section 7.3.

7.2.3 Variation 1

This deployment scheme, using a minimal set of centralized PKI servers, does not produce an overall bandwidth requirement much different than Option 1. The model for Variation 1 was only run at the five times traffic level, and the total bandwidth used is 7.72 Mbps, as compared to 7.714 Mbps for Option 1 with the same traffic multiplier.

An area of concern with this configuration is individual link utilization. Table 7-3 shows the same links that were presented in Table 7-2. As expected, there are only minor differences between the tables seen here. Exceptions are the Masset and the Trenton-Tunney's Pasture links. The Trenton-Tunney's Pasture link increased in utilization from 25.11 % to

35.73 %. This is not as significant as the increase for the Masset-Borden link, which jumped from 44.14 % to 67.53 %. This indicates that the Masset-Borden link should be upgraded from its current 56 Kbps capacity. Also notice that links 66 and 65 remain at their baseline traffic load levels, indicating that the PKI traffic doesn't use these links.

Table 7-3 Variation 1 Link Utilization

Link #	Link Capacity	Link Source	Link Destination	Link Utilization
4	256 Kbps	Tunney's BAR 1	Wainwright	42.75%
12	256 Kbps	Tunney's BAR 1	Suffield	31.81%
12	256 Kbps	Suffield	Tunney's BAR 1	24.06%
6	384 Kbps	Tunney's BAR 1	Shilo	30.06%
14	56 Kbps	Tunney's BAR 1	Yellowknife	26.49%
21	56 Kbps	Masset	Borden BAR 1	67.53%
22	56 Kbps	Masset	Tunney's BAR 3	6.8%
30	384 Kbps	Trenton	Tunney's BAR 2	35.73%
30	384 Kbps	Tunney's BAR 2	Trenton	33.45%
56	768 Kbps	Tunney's BAR 1	Gagetown	25.9%
66	1.536 Mbps	Borden BAR 3	Tunney's BAR 3	50.68%
66	1.536 Mbps	Tunney's BAR 3	Borden BAR 3	28.92%
65	1.536 Mbps	Pearkes BAR	Borden BAR 1	31.61%
65	1.536 Mbps	Borden BAR 1	Pearkes BAR	43.92%

Variation 1 will incur the same bandwidth costs as the Option 1 configuration using the five times traffic multiplier. The hardware and personnel costs of this scheme are explored in section 7.3.

7.2.4 Variation 2

This deployment, using centralized CAs and distributed DSAs and TSAs has a bandwidth requirement similar to the baseline model. Variation 2 used a total of 6.123 Mbps of the available bandwidth, which includes the baseline traffic and the PKI traffic. Recall that the baseline used 6.11 Mbps of bandwidth. As was determined for Option 2 and 3, there is an inconsequential difference in bandwidth used.

The ISAP CPU utilization is similar to all of the other models, and can be considered well within the capabilities of the Cisco 7500 routers, and therefore of no consequence in a cost calculation.

7.3 Cost Implications

This cost analysis section looks at the personnel and equipment requirements needed to maintain and run each of the PKI deployment configurations described in this thesis. Before this can be done, some basic cost concepts must be presented.

Each deployment will have an annual recurring cost, and an initial deployment cost. The annual recurring costs include the price of the DWAN backbone bandwidth utilization,

and the personnel costs. Initial deployment costs include the necessary hardware, software, and personnel to integrate the PKI with the DWAN.

Within the DND, the association of a dollar value with the cost of employing a person is termed a Person-Year (PY). A PY is a value that includes the salary and benefits that a person qualified for a position is entitled to. Benefits include medical and dental coverage, for example. An assumed value for a PY is \$80,000, and this value will be used throughout this section for personnel cost calculations.

The equipment that is needed for each PKI configuration includes the servers that host the CAs, DSAs, and TSAs. Other equipment vital to the PKI is the equipment providing a secure and trusted source of time for the TSA servers. Each of the PKI service applications can be installed on Commercial Off-The-Shelf (COTS) computing platforms costing no more than \$5,000 each, and that figure is used for each required computer. This platform total includes the costs for the computer itself, network connection, and system integration costs. The personnel that will administer the PKI, locally and remotely, will need some training, and this cost is also included in the \$5000.

The secure time source, on the other hand, can vary in price from less than \$100 for a modem, to over \$12,000 for a high end GPS receiver system. For the purposes of this thesis, the assumption is made that a GPS receiver system such as the CNS Clock (CNS Systems Inc) will be the trusted source of time for the TSAs. This particular system retails for \$1,400, and it can supply a time signal to two separate servers simultaneously.

Each of the deployment configurations will need some redundancy and testing/prototyping capability, and five PKI servers should be included for these reasons. This adds \$25,000 to the cost of each scheme.

Now that each PKI component and the personnel needed to administer the system have been assigned a cost, the total cost implications for the PKI deployment schemes can be presented.

7.3.1 Option 1 Costs

The Option 1 deployment scheme, which is the current DND deployment plan, uses 15 servers to host the PKI software. Each of the CAs, DSAs and TSA is running on a separate platform, and this would cost \$75,000 based on the \$5,000/platform figure. There would be a need for backup servers and test-bed systems as well (five), bringing the total to \$100,000 for the server platforms.

The TSA time source is required. The DND would need at least two systems; one being the backup. This adds \$2,800 to the total cost of the deployment.

Personnel costs will be the largest expenditure for any of the PKI configurations. The current DND plan calls for 5 personnel to administer the central PKI servers. This plan also calls for 1 PKI Administrator for every 100 users at the base level (LRAs). This will not actually occur, since a large base with 2000 personnel would not use 20 PKI Administrators. Since there must be at least three personnel at each base whose primary or secondary function is to be a PKI Administrator, each base is assigned a number of PKI personnel based on their relative populations. Three personnel are the minimum, because it takes two to recover a user profile, as dictated by the Certificate Policy for DND. The third person is needed to cover the times when one of the other two is unavailable. Each base is assigned a ranking of small, medium, or large with respect to their number of potential PKI users. The number of PKI

Administrators is related to these size rankings. Table 7-4 shows how the bases are ranked according to size, and the number of personnel need to administer the PKI at the CF Base level.

Table 7-4 CF Base PKI Personnel

Size	Bases	PKI Personnel
Small	Wainwright, Suffield, Yellowknife, Aldergrove, Masset, North Bay, Gander, St. John's, Goose Bay, Millcove	3
Medium	Shilo, Moose Jaw, Comox, Petawawa, Bagotville, Toronto, Shearwater, London, St. Jean	4
Large	Winnipeg, Edmonton, Esquimalt, Cold Lake, Kingston, Trenton, Greenwood, Halifax, Gagetown, Valcartier, Longue Pointe, Borden, NDHQ	6

The number of personnel derived from Table 7-4 is 144. It is not reasonable to assume that this number of personnel will be doing strictly PKI administrative duties all of the time. After everyone on a CF Base has been issued their PKI credentials and are able to use the PKI properly, the workload for the PKI Administrators will be reduced. This end-state posture is assumed to have been reached for the cost calculations. It is assumed that the same number of personnel will perform PKI duties as listed in Table 7-4, and that they will be assigned this function as 25% of their duties. This identifies a personnel requirement of 36 PYs for the CF Bases, plus the 5 PYs required for the CA managers previously identified. Therefore, 41 PYs

are required for the Option 1 PKI deployment. At \$80,000/PY, the annual cost for the personnel of Option 1 is \$3,280,000.

Recall that the bandwidth needed for the Option 1 model with the five times traffic multiplier was \$12,000. All costs are now added up in Table 7-5. There are two cost columns; one for the deployment cost, and one for the annual recurring cost.

Table 7-5 Option 1 Total Costs

PKI Need	Deployment Cost	Recurring Cost
Bandwidth		\$12,000
Platforms	\$100,000	
Time Source Equipment	\$2800	
Personnel		\$3,280,000
Totals	\$102,800	\$3,292,000

7.3.2 Option 2 Costs

The Option 2 deployment configuration started with the Option 1 scheme and added a third level of CAs, DSAs, and TSAs at the base level. The bandwidth costs described in section 7.2.2 for Options 2 and 3 are negligible and not included in the total calculations for Options 2 and 3.

In Option 2, there is an additional CA, DSA, and TSA at each of the 31 CF Bases not located in the Ottawa area. The level 2 PKI servers located in Ottawa serve the PKI requirements of the International, NDHQ, and Pearkes sub-nets. Therefore, 93 additional servers are required to deploy Option 2, costing \$465,000 more than the server cost of Option 1. The total for the PKI platforms is \$565,000. There will also be a need for 31 additional GPS time source systems, costing \$43,400. The two centralized time source systems, costing \$2,800, are still required, bringing the total to \$46,200 for this equipment.

With each base managing its own CA, there must be a full-time CA administrator in addition to the personnel requirements detailed for the Option 1. This adds 31 PYs to the total personnel requirements for the Option 2 PKI. This brings the total personnel requirement for the PKI up from the 41 needed for Option 1 to a total of 72 for Option 2. The totals for Option 2 are provided in Table 7-6.

Table 7-6 Option 2 Total Costs

PKI Need	Deployment Cost	Recurring Cost
Bandwidth		negligible
Platforms	\$565,000	
Time Source Equipment	\$46,200	
Personnel		\$5,760,000
Totals	\$611,200	\$5,760,000

The recurring cost of Option 2 is almost double that of Option 1, but the specific benefits and disadvantages may make it a feasible deployment scheme, and these features are presented in section 7.5.

7.3.3 Option 3 Costs

The difference between the Option 3 and Option 2 deployments is the removal of the centralized level 2 PKI servers. This scheme will reduce the number of personnel (5 PYs) needed at the central locations, but will still require a significant number of personnel at the base level. Since there is a need for at least three CA Administrators in the Ottawa area, the number of personnel needed in Ottawa for the PKI is reduced from a requirement of 5 PYs to 3 PYs. These three personnel would have the responsibility to maintain the region's CA and monitoring the overall health of the PKI. The total number of PYs needed to operate the Option 3 deployment is 70, where 36 are PKI Administrators at the bases, 31 are CA Administrators, and 3 are CA Administrators at Tunney's Pasture.

Option 3 eliminates the need for 10 of the centralized servers in Borden and Ottawa. This reduces the PKI server costs from \$565,000 for Option 2 to a total of \$515,000. The time source hardware is still required; 31 systems for the bases and the 2 centralized systems at Tunney's pasture, for a total of 33 at the cost of \$1,400 each.

The totals for this configuration are depicted in Table 7-7.

Table 7-7 Option 3 Total Costs

PKI Need	Deployment Cost	Recurring Cost
Bandwidth		negligible
Platforms	\$515,000	
Time Source Equipment	\$46,200	
Personnel		\$5,600,000
Totals	\$561,200	\$5,600,000

7.3.4 Variation 1 Deployment Costs

This deployment simplifies the centralized scheme of Option 1 by reducing the number of PKI servers to the bare minimum. The scheme uses 3 CAs and DSAs, and 2 TSA servers, for a total of 8. Adding the 5 for redundancy and testing would produce a total cost figure of \$65,000 for the PKI servers. Two time source systems are used in this deployment, plus the spare system at Tunney's Pasture, costing \$4,200.

The personnel needed to maintain and administer this deployment scheme are assumed to be similar to the personnel needs of Option 1. The only difference with this variation is that 1 more person would be needed since the central servers are in two different physical locations, CFB Borden and Tunney's Pasture. The CAs, DSAs, and TSAs can be administered by three full-time positions at each of the locations, for a total of 6 PYs. The 36 LRA personnel at the bases are still required. The total for personnel is 42 PYs, at a cost of \$3,360,000 per year.

The bandwidth needs of this deployment are the same as those of Option 1, costing \$12,000 annually.

Considering the long term cost of the extra CA Administrator PY used in this variation, installing the western region CA at Tunney's Pasture would reduce the overall PY requirement by one. This is a real savings, when the life-cycle of the PKI is considered. Assuming the equipment is suitable for a seven year deployment, moving the western CA to Tunney's Pasture would save \$560,000, calculated as the cost of a PY for seven years. The DSA and TSA for the western CF bases remain at CFB Borden, and there would be no significant change in any link utilization.

Table 7-8 summarizes the total system costs for Variation 1.

Table 7-8 Variation 1 Total Costs

PKI Need	Deployment Cost	Recurring Cost
Bandwidth		\$12,000
Platforms	\$65,000	
Time Source Equipment	\$4,200	
Personnel		\$3,360,000
Totals	\$69,200	\$3,372,000

7.3.5 Variation 2 Costs

This final deployment configuration uses centralized CAs, and distributed DSAs and TSAs. The personnel costs of Variation 2 are calculated from the requirement to have 36 PYs used for PKI Administration at the base level, and 6 PYs for CA Administration at the two central locations. The total requirement of 42 PYs will cost \$3,360,000.

The number of required PKI servers would be less at the Borden and Tunney's Pasture sites since the DSAs and TSAs would all be located at the bases. The Master DSA and the NDHQ DSA are still needed at Tunney's Pasture in Ottawa. The CAs are still centrally located, but only 3 are used; a level 1 and a level 2 CA at Tunney's Pasture, and another level 2 at CFB Borden. A master TSA is located at the Tunney's Pasture site. This brings the total number for the required server platforms to 62 for the CF Bases (2 each), 5 at Tunney's Pasture, and the CA server at CFB Borden. The five additional testing and redundancy platforms bring the total to 73, worth \$365,000.

The calculation for the time source systems cost is based on the assumption of individual CNS Clocks for all 33 TSAs (31 at the bases, 2 at Tunney's Pasture), totaling \$46,200.

The total cost for the Variation 2 deployment is shown in Table 7-9.

Table 7-9 Variation 2 Total Costs

PKI Need	Deployment Cost	Recurring Cost
Bandwidth		negligible
Platforms	\$365,000	
Time Source Equipment	\$46,200	
Personnel		\$3,360,000
Totals	\$411,200	\$3,360,000

To reduce the long term costs of Variation 2, the CA at CFB Borden can be installed at Tunney's Pasture, and a savings of \$560,000 over seven years can be realized.

7.3.6 Costs Summary

Looking at the lifecycle cost for the various deployment schemes identifies a few points of interest. First, the human PY costs are the major expense for every deployment option and variation. The equipment communications infrastructure costs are relatively low. The platforms and time source equipment should be usable for at least seven years. After this time, the equipment may need to be upgraded.

One of the costs not examined in detail for the PKI system is the cost of the required software. The DND has an 'unlimited' site license for all of the necessary Entrust PKI software packages. This means that this cost has already been incurred.

Table 7-10 shows the total deployment and recurring annual costs of the five deployment schemes using the five times traffic multiplier. The life cycle cost column represents the totals of the deployment cost and seven years of recurring costs.

Table 7-10 Deployment Costs Summary

PKI Deployment Scheme	Deployment Cost [Thousands]	Annual Recurring Cost [Millions]	Life Cycle Costs (7 Years) [Millions]
Option 1	\$103	\$3.3	\$23.2
Option 2	\$611	\$5.76	\$40.1
Option 3	\$561	\$5.6	\$39.8
Variation 1	\$69	\$3.38	\$23.7
Variation 2	\$411	\$3.36	\$23.9

7.4 PKI Security Analysis

This section examines the security strengths and weaknesses of each deployment scheme. Before this is done for each individual configuration, the risks to the PKI are reviewed.

Each CF Base has at least one physically secure facility where sensitive computing components can be installed. These facilities have restricted personnel access lists and provide the necessary physical mechanisms to prevent an intruder from gaining access to the equipment. A security policy detailing the security measures for PKI hardware, software, and the physical environment of a CA system must be in place to protect the PKI. Within DND, PKI security policy and procedures are established and administered by the Level 1 CA. This CA will automatically record every event related to the PKI and host operating system (OS) management functions. The DND PKI security officers will process the security audit logs weekly, reviewing each entry and taking any necessary actions to resolve security incidents (14).

Access to the OS and PKI software installed on each platform must be limited to the few personnel who need to administer the systems. Each platform should only be running the required OS and the specific PKI software (either CA, DSA, or TSA), and nothing else. The access control features of the OS must be configured and maintained to meet the security policy.

The DWAN is susceptible to an external attack because of its connection to the Internet through the Interim Firewall Service (IFS). This firewall strictly controls the acceptable external IP addresses and provides some limited protection from the insecure Internet. The DWAN uses leased lines from a commercial provider, and the risk of an attack on these lines has to be considered. An attacker could gain control of the DWAN backbone lines or switches and either deny service or alter the system configuration in such a way as to cause a severe degradation in quality of service. Traffic could be intercepted and compromised if messages are not protected using encryption techniques (e.g. PKI). These are general DWAN security issues that exist independently of the deployment of the PKI, although the PKI may be used as an enabling technology to mitigate some of the risks. The focus of the following discussion is on the PKI servers themselves.

The PKI server with the highest risk factor is the CA. A compromise of a CA would allow an attacker to seriously jeopardize the entire PKI's integrity. Legitimate users can be removed, false PKI credentials can be created, and cross-certifications can be removed or created. Other dangers to the PKI resulting from a CA compromise include information disclosure, denial of service, the changing of operational parameters, and the deletion or alteration of critical system files. Some of these files include the CRLs, archived keys, and the system initialization files that identify the location of the other PKI servers. The greatest risk to

a CA is from within the DWAN system perimeter. Reducing the number of CAs and the personnel with access to them will reduce the administrative complexity and enhance the overall infrastructure security.

The publicly available DSA servers have a much lower risk factor than the CA servers. The compromise of a single DSA platform, other than the Master DSA, allows an attacker to deny the services of the particular server. The certificates of each DSA are duplicated on every other DSA, and a user can receive a proper certificate from an alternate site without even knowing which DSA actually provided the service. The compromise of the Master DSA would have more serious consequences. Every other DSA shadows the Master and contains copies of the certificates provided by the Master DSA. Assuming that a CA hasn't been compromised, an attacker with control of the Master DSA can stop the transfer of valid certificates and CRLs to each of the shadow DSAs.

If a TSA server is compromised, full non-repudiation may become unavailable. If only a single TSA platform is unavailable, a user may get time-stamps from a secondary TSA server. If the TSA system clock is questionable due to the removal of the time source equipment for that server, or there is a successful penetration by an attacker, false time-stamps could be generated. The alteration of the TSA system time is detectable by PKI security officers, as TSA events are also logged for audit purposes.

With all of the described security threats, each PKI deployment scheme can be analyzed to determine the risk to each particular configuration. The level of security associated with a PKI deployment is only discussed in general. The actual security of the system is, as always, determined by a complete security environment comprising physical, procedural, personnel, communications, and computer security.

7.4.1 Option 1 Security Analysis

The centralized Option 1 deployment has all of the CA servers installed at a physically secure facility at Tunney's Pasture in Ottawa. This feature provides the highest possible security for these sensitive components. Access to the facility is limited and access to the PKI servers is even more so. Only those personnel with the need to have access to the hardware can enter the actual room where the PKI servers are located. Of these personnel, only those with the need to perform administrative functions have the necessary rights to login to the CA servers. This centralized control provides strong assurances of restricted access control for the CAs. This mitigates the risk of CA compromise.

The Master DSA and TSA are located in the same facility in Ottawa, minimizing the threat to these servers as well. The other less critical DSAs are located at CFB Borden in the secure computing facility found there. A possible improvement that might be made to the security posture of this deployment scheme would be the transfer of the DSAs located at CFB Borden to the facility in Tunney's Pasture. This would force some of the PKI message traffic on the DWAN to use the already busy link between Borden and Ottawa.

Overall, Option 1 provides a very secure environment for the PKI assets. The LRAs located at each base can only remotely access a CA for the specific functions of creating, recovering, or removing a user profile.

7.4.2 Options 2 and 3 Security Analysis

These two schemes are very similar from a security standpoint, and are examined together in this section. The centrally located servers of these deployments have the same security as Option 1. The main threat to these deployments is to the distributed PKI servers, especially the CAs.

The possible CA compromise scenarios described in the beginning of section 7.4 represent the highest risk to the PKI integrity for these two schemes. The TSAs and DSAs are also more vulnerable, since there are so many more of them in the infrastructure.

The increased number of personnel who must have CA privileges, and the associated complexity of administration, presents a real danger to the PKI. Only one CA needs to be compromised before the entire PKI integrity becomes suspect. Each base has its own secure facility for the PKI servers, but the higher number of attack points represents a real threat.

The higher security vulnerability of these two deployment options does not exclude them from deployment. There has to exist some amount of trust in the personnel identified as PKI administrators, or the PKI would never be deployed at all. The risk of an attacker exploiting the system vulnerability may be mitigated by procedural security controls. Options 2 and 3's main advantage of a low DWAN backbone link utilization may offset the effort of establishing the added security controls.

7.4.3 Variation 1 Security Analysis

This deployment in theory provides more assurance of security than Option 1. There are fewer PKI servers in use, and they are all at centrally located secure facilities. The number of personnel needed to be CA administrators is slightly reduced from Option 1, and this fact further reduces the chance of a compromise. This deployment scheme is considered to have the highest security assurance of all options.

7.4.4 Variation 2 Security Analysis

The Variation 2 deployment has the CAs centrally located, and therefore the most sensitive PKI servers would have the maximum protection.

The main threat to this scheme is against the distributed DSAs and TSAs. The scenarios described in the first part of section 7.4 for the DSAs and TSAs indicated that a compromise of any of these servers would be less damaging than a CA compromise. A straight-forward denial-of-service for either of these components could actually be transparent to most users. There would be a slightly longer processing delay, as the needed certificates or time-stamps would have to come from another location on the DWAN.

The overall threat to Variation 2 can be considered much lower than Options 2 and 3, but higher than Variation 1 and Option 1.

7.5 Recommended Deployment Scheme

Careful consideration of the performance, cost, complexity, and security characteristics of the tested PKI deployments leads to the recommendation of a single scheme. Table 7-11 provides a summary of the deployment schemes and their strengths and weaknesses.

Table 7-11 PKI Deployment Summary

Configuration	High Bandwidth Use	High Cost	Complex	CA Security
Option 1	Yes	No	No	Yes
Option 2	No	Yes	Yes	No
Option 3	No	Yes	Yes	No
Variation 1	Yes	No	No	Yes
Variation 2	No	No	No	Yes

Option 1 has the advantage of a centrally controlled and securable deployment. The bandwidth utilization for the PKI messages is high, but not unreasonable, given the available bandwidth during peak hours. The lifecycle cost is also an advantage, as it is considerably lower than Options 2 and 3.

Options 2 and 3 have the highest lifecycle costs associated with their deployments, due to the need for so many personnel to administer the PKI. The security factor is also a disadvantage of these two schemes, presenting many more points of attack to the PKI, and particularly the CAs. The small bandwidth utilization increase represents what can be considered the only advantage of Options 2 and 3.

Variation 1 has all of the same features as Option 1, and a similar lifecycle expenditure is required to deploy and administer the system. This scheme should be considered as an alternative to Option 1, as it reduces the system complexity somewhat. The security is slightly better than Option 1, due to the reduced number of servers and personnel needed for the PKI.

Finally, Variation 2 presents a viable compromise between all of the other described deployments. The lifecycle cost is comparable to Option 1 and Variation 1, the bandwidth requirements are low, and the CA servers are securable.

As the available DWAN bandwidth is reduced because of the introduction of more applications, the bandwidth requirements of the PKI will become a more important factor. Variation 2 provides local access to certificates, CRLs, and time-stamps, thereby virtually eliminating an increase in DWAN backbone link utilization and processing delays.

For the needs of the DND, the Variation 2 PKI deployment is the recommended configuration.

8 DISCUSSION AND CONCLUSION

8.1 Introduction

This chapter presents several discussion items relevant to the DWAN and PKI deployment. The importance of PKI deployment research as it applies to the DND is presented here. Points discussed include PKI deployment options that are suited to the current DWAN configuration, and PKI deployment options that may become a better choice as the DWAN evolves. The most advantageous use of the available Entrust PKI options, including file compression, time-stamp settings, and efficient cryptographic algorithms, is included in this chapter.

The DWAN models developed for this thesis can be used in other areas of research. For example, these models can be used for testing the impact of implementing new applications on the DWAN and for fault analysis testing, including single or multiple fault scenarios.

After the discussion items have been explored, areas for further research on PKIs are introduced, followed by the conclusion.

8.2 Discussion

The primary objective of this thesis was to explore the various PKI deployment options available to the DND. As a result of the research into the most likely types of deployment schemes, and how they will affect the DWAN performance, the PKI can be

deployed with a greater level of confidence. The scheme that best matches the DND PKI Policy objectives and cost constraints can be chosen for implementation.

8.2.1 DWAN Baseline Model

The baseline DWAN model can be used to test the integration of any new application, provided the performance characteristics are known. These characteristics include protocols, hardware configurations, bandwidth needs, and any other information that will make the resulting model as realistic as possible.

Within the next year, several DND applications will be migrated to the DWAN from other networks. One of these applications is the MMHS, the military message system. Others include the command and control applications for the CF air and land forces. These systems will carry mission critical messages that are needed for the efficient preparation, deployment, and control of various CF formations. Using the DWAN baseline model, each separate application can be quickly and safely tested in a realistic environment. The units at NDHQ responsible for the successful integration of applications with the DWAN have expressed their interest in using the specified DWAN baseline model for initial testing purposes.

There are currently plans in the works that will see the bandwidth increased across the entire backbone, including the links between the BARs at CFB Borden and Tunney's Pasture. The DWAN baseline model can be easily updated to evolve along with the network it emulates. The DWAN traffic volumetrics, which are continuously being collected by NDHQ, provide the means to update the DWAN baseline model with the most recent and accurate traffic parameters.

8.2.2 DWAN PKI Models

The DWAN PKI models specified in this thesis were developed to provide realistic options for a successful PKI deployment. The results described in chapter 7 will allow the NDHQ DDCEI personnel to proceed confidently with the DND PKI deployment. The deployment option that best meets DDCEI's mandate can now be implemented without the DDCEI personnel performing time-consuming alternative deployment testing.

The Variation 2 deployment has been recommended as the most suitable scheme for the DND PKI given the current DWAN capabilities. This deployment offers the best security posture for the CA servers while distributing the high traffic DSA and TSA servers to each CF base. As the DWAN links experience utilization increases due to the integration of new applications and an increase in Internet capability, the Variation 2 scheme offers the best compromise between security, cost, and link utilization. Each of the DSAs and TSAs can be administered locally by the PKI Administrators at each CF base. The TSA servers are stand-alone, and they rely on a local system for their time source. The DSAs mirror the Master DSA at Tunney's Pasture, and updates to the CF base DSAs can be scheduled to coincide with off-hours. CRL updates are also schedulable for off-hours, using the DWAN backbone links when they have the least utilization.

As the DWAN continues to evolve to incorporate technological advances, and the DWAN baseline is updated to reflect these changes, the DWAN PKI models will still be of value. For example, when the capacity of the backbone links is increased to levels high enough to make a completely centralized PKI deployment more desirable, the Option 1 and Variation 1 deployment schemes can be tested again on an updated DWAN baseline model.

All of the specified models can be used to perform fault testing on a specified network. For example, the current plan for the DND PKI deployment has most of the PKI servers connected to Tunney's BAR 1. What if this Cisco 7500 router suffers a disabling failure? How quickly can a spare component be brought on-line? What will happen to each backbone link utilization if this component fails? Each of these questions can be answered using the DWAN/PKI models. ComNet allows the user to fail a component or link at any time during a simulation, and the results can be measured. This feature of ComNet will allow the network engineers to test any possible failure scenario on the DWAN, and help them plan how to avoid, or minimize, the impact of component and link failures.

8.2.3 Issues Relating to Entrust PKI Options

The Entrust PKI software applications allow the selection of several options that will directly affect the performance of the PKI deployment. Each of the symmetric key encryption/decryption ciphers that are available within Entrust has a performance cost associated with it. While DES has a relatively high throughput compared with the other available block ciphers, it is not considered secure enough any longer. The CAST algorithms, which have been approved for use for the Government of Canada by the CSE, have a slower throughput than DES, but are considered secure. For digital signatures, the RSA public-key scheme is the algorithm of choice. Entrust also allows the PKI administrator to select the Digital Signature Algorithm (DSA), if desired. RSA is also used for the encryption of the symmetric keys (CAST) used for encryption and decryption within the DND PKI.

The Entrust TimeStamp subsystem can also be configured with the client software, Entrust Entelligence. The user can select when a timestamp will be applied; for example, whenever something is encrypted and/or signed. This is an important feature that may be used inappropriately, causing a cost increase within the DND for the use of this PKI subsystem. If the DND PKI Policy is expanded to include a requirement to digitally sign all e-mail messages, and the timestamp feature is selected for signing, timestamps will be used for every e-mail, however inconsequential they are. Currently timestamps are sold with a per/unit cost, as a notary service. Timestamp capacity is purchased in blocks, and when consumed, additional capacity must be purchased from Entrust. This is the case even though DND is hosting the timestamp servers. If timestamps are used with everyday unclassified e-mail, they would be wasted, not to mention the increase in traffic on the networks resulting from the timestamping function. The Government of Canada is currently engaged in the process of purchasing unlimited timestamps for the GOC PKIs, as part of an annual procurement contract for the Entrust PKI systems.

Another option available with the Entrust PKI is the use of file compression. Assuming most messages using the PKI are text-based, file compression can significantly reduce the bandwidth requirements of the entire DWAN backbone. The compression is applied to the document before it has been encrypted, signed, and/or timestamped. Recalling that one of the assumptions made for the PKI models was that a message would have been sent with or without a PKI, file compression would reduce the overall bandwidth used in every model examined in this thesis. Depending on how many applications use the PKI, the bandwidth needs for some PKI deployment options could be less than the baseline model, if file compression is used.

8.2.4 DND PKI Certificate Use

The DND PKI Policy outlines the general rules for the Distinguished Names (DNs) used in the X.509 certificates. It does not state exactly how individuals will be assigned these DNs so that uniqueness is guaranteed throughout the DND.

The use of surnames and initials with a CF rank does not ensure a unique DN will be formed. Another attribute, such as the last three numbers of a CF Service Number (SN), must be included in the DN to guarantee uniqueness.

A certificate can also be issued against a position title instead of an individual. In these cases, the selection of a DN can be difficult, since the same position title can be used at different CF bases. Various position acronyms can have a completely different meaning in another organization, and subsequently, caution must be exercised when using a DN that is position based.

8.3 Further Research Areas

Research in the computer security field continues to be a hot topic in the IT world. Recent attacks on several public e-commerce sites have prompted a scramble within the business world to put computer security measures into place as quickly as possible.

The successful deployment of a PKI is an enabling technology that will protect an organization's vital information, and provide the non-repudiation services needed to protect the integrity of contractual agreements that use digital signatures and timestamps.

One area of research that has not yet produced a comprehensive solution involves PKI policy compatibility. The DND PKI Policy and CPS contain all of the details about how and when PKI certificates will be used, and by whom.

The mechanism to cross-certify CA's is available, but how will the individual PKI policies be mapped to each other? The process of policy compatibility checking may be automated, to a certain extent, but not completely. The use of an organization's legal services will be required to help with the decision of cross-certification outside of a CA's domain. Object Identifiers (OIDs) within a certificate can be used to indicate whether or not another CA's certificate policy is acceptable, and for what specific or general purpose. For example, the DND PKI may be cross-certified with a contractor to enable secure communications and transactions between a specific department within DND and the contractor. Once the legal departments of both parties have agreed that the CAs can be cross-certified, and identified the specific individuals and/or rights this cross-certification applies to, the necessary certificate OIDs can be included. This use of an OID in this instance is still not clearly defined by the X.509 certificate standards, and is still a discussion item for the IETF PKIX working group.

Another research area involving PKIs deals with the use of the X.509 certificates without a public-key. The primary purpose of a public key certificate is to bind an identity with a public key. What if the same certificate is used to bind an identity with a set of attributes or privileges instead of a public key? This type of certificate, called an attribute certificate (AC), is used in a Privilege Management Infrastructure (PMI). PMI certificates differ from PKI certificates in that they provide a mechanism for an organization to grant privileges to an individual for use within that organization.

As an example of how this infrastructure works, the following scenario provides an example of one possible use. A CF aircraft technician has certain qualifications and authorizations that allow him/her to sign for work performed on an aircraft. These qualifications and authorizations are mandatory to ensure aircraft airworthiness requirements are met. The attainment of an authorization means that a technician has taken the necessary training and has acquired enough experience to do the work.

An X.509 attribute certificate can capture these qualifications and authorizations, so that when a technician tries to digitally sign for work completed, the applicable attribute certificate can be accessed to ensure the technician has the necessary authorization to sign for this work. If he/she is authorized according to the PMI, the technician signs for the work using the normal PKI functions.

8.4 Conclusion

The successful deployment of a PKI on a WAN requires a careful study of the nature of the target network, and an understanding of the interaction between PKI components. A simple centralized installation could produce a significant performance degradation, especially on a WAN that already has high link utilization due to the type of applications in use, the nature of everyday business, or Internet traffic.

A PKI provides essential security services including data integrity, non-repudiation, digital signatures, and data security. More importantly, a PKI associates a unique entity with their public key in a standard and exchangeable certificate, and automatically manages key lifetimes and archiving based on the PKI policy.

This thesis has introduced several possible PKI configuration alternatives and explained how the results of each decision can be captured using realistic models of the DWAN. The problem of a variable size population at each base was resolved using the known traffic ratio available from the MMHS message system. While it is possible to drastically reduce WAN PKI traffic using a distributed deployment scheme with the CAs, DSAs, and TSA co-located at the CF bases, the resulting system complexity, personnel and equipment costs, and PKI security weaknesses could prohibit such a selection.

A likely compromise for the CF is the use of centralized CAs, with the DSAs and TSAs dispersed to the bases and units. DWAN traffic increases are minimal, while the central system management concept reduces costs significantly and enhances the security of the CAs.

The DND has several PKI deployment options available to choose from, and the impact of using each configuration is documented in this work. A deployment scheme that is suited for today's DWAN characteristics may need to be reconsidered as the DWAN evolves with technological advances.

Several operational models are now available for further network studies, and each model can be rapidly updated to continuously mirror actual DWAN characteristics. The models provide the means to test new applications and configurations, without incurring the expense and risk of performing trial and error installations on the actual system.

LIST OF REFERENCES

1. Entrust. *What is a PKI?* Entrust Technologies White Paper, 1999
2. J. Lewis. *Public Key Infrastructure Architecture*. The Burton Group Network Strategy Report, v1. 9 July 1997
3. D.R. Stinson. *Cryptography: Theory and Practice*, CRC Press LLC, 1995
4. W. Diffie, M.E. Hellman. *New Directions in Cryptography*. IEEE Transactions on Information Theory, 22 , pp 644-654, 1976
5. R.L. Rivest, A. Shamir, L. Adleman. *A Method for Obtaining Digital Signatures and Public Key Cryptosystems*. Communications of the ACM, 21, pp 120-126, 1978
6. W.J. Caelli, E.P. Dawson, S.A. Rea. *PKI, Elliptic Curve Cryptography, and Digital Signatures*. Computers & Security, Vol 18 No 1 1999
7. C. Adams, S. Lloyd, *Understanding Public-Key Infrastructure*, Macmillan Technical Publishing, 1999
8. S. Haber, W.S. Stornetta. *How to Time-Stamp a Digital Document*. Journal of Cryptology, Vol 3, 1991
9. D.L. Mills. *Network Time Protocol v. 3 Specification, Implementation and Analysis*. RFC 1305, Network Working Group, March 1992
10. D.E.R. Denning. *Cryptography and Data Security*. Pp 169-170, Addison-Wesley Publishing, 1983
11. S. Kent. *Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management*. RFC 1422 Network Working Group (IETF), February 1993
12. Treasury Board of Canada, *Secretariat Policy for Public Key Infrastructure Management in the Government of Canada v 1.0*, 29 January 1999
13. DDCEI 3-5-ETL. *Concept of Operations for the National Defence Public Key Infrastructure v. 2.0*. 15 March 1999
14. DND CA. *Certification Practice Statement v. 0.3*. May 1999.

15. C. Adams, *Constructing Symmetric Ciphers Using the CAST Design Procedure*, Selected Areas in Cryptography, pp71-104, vol12, no. 3, November 1997
16. Cisco Systems Inc. *Synchronizing Clocks with NTP Service*. Cisco Systems Inc, 1997
17. 4th CACR Information Security Workshop, Third-Party Evaluation Criteria, Centre for Applied Cryptographic Research (CACR), University of Waterloo, November 8, 1999
Pitney Bowes, Corporate Headquarters, Stamford, Connecticut.
18. DTSES 4-5/PM DWAN ECP, *Project Implementation Plan*, DWAN Extension and Consolidated Project, June 1999
19. DIOM 3-5-4-2, *National Systems Overview*, Security Management – National ISSO (DWAN), 6 Oct 99
20. Beaton, A. Olmo, N. Siddiqui. *Virtual Private Networks: Here or Hype?*. kpmg-vpn.se-com.com white paper, 1998
21. NIST . *Cryptography Standards and Infrastructures for the Twenty-first Century*. <http://nist.se-com.com> 1998
22. A.M. Law, W.D. Kelton, *Simulation Modeling & Analysis*, pp 307 – 314, McGraw-Hill, 1991
23. DTCEI 4-2-7, *DWAN Architecture Document*,
<http://131.134.0.39/cfrets/TISO/dwanarch.htm> , Apr 1997
24. Tanenbaum, Andrew S., *Computer Networks*, Prentice Hall, 1996

BIBLIOGRAPHY

1. B.-J. Koops. *Crypto Law Survey - Canada*. <http://cwis.kub.nl/~frw/people/koops/lawsurvey.htm>. July 1999.
2. J. Postel. *DARPA Internet Program Protocol Specification*. RFC 791 Information Sciences Institute USC Marina del Rey, September 1981.
3. Entrust. *Performance Benchmarks*. Entrust Technologies Limited 1998
4. S. Chokhani, W. Ford. *Certificate Policy and Certification Practices Framework*. RFC 2527 The Internet Society, March 1999.
5. RSA Labs FAQs. <http://www.rsa.com/rsalabs/faq>
6. R.A. Clyde. *Enterprise Security: Built on Sound Policies*. <http://axent.se-com.com> Montgomery Research
7. D.H. Crocker. *Standard for the Format of ARPA Internet Text Messages*. RFC 822 Dept of Elec Eng, U of Delaware, 13 August 1982.
8. N. Asokan, G. Tsudik, M. Waidner. *Server-supported signatures*. Journal of Computer Security 5 1997
9. RSA Security Protocols Overview. <http://rsa.com/standards/protocols>
10. M. Winslett, N. Ching, V. Jones, I. Slepchin. *Using Digital Credentials on the WWW*. Journal of Computer Security 5 1997
11. M.J. Bell . *Digital Signatures: The Key to Information Technology Security*. <http://aba.se-com.com> 1998
12. B.S. Kaliski Jr. *Examples of the PKCS Standards*. RSA Laboratories Technical Note, 1 November 1993
13. T. Patterson. *Establishing Electronic Commerce via Certificate Authorities*. Cyberguard-ecom.se-com.com white paper, 1997
14. J. Hughes. *Certificate Inter-Operability White Paper*. Computers & Security Vol 18, No. 3, 1999

15. Entrust. *PKI Features and Benefits*. Entrust Technologies 1999
16. R. Housley, W. Ford, W. Polk, D. Solo. *Internet X.509 PKI Certificate and CRL Profile*. RFC 2459, Network Working Group, The Internet Society, January 1999
17. C. Adams, S. Farrell. *Internet X.509 PKI Certificate Management Protocols*. RFC 2510, Network Working Group, The Internet Society, March 1999
18. M. Myers, C. Adams, D. Solo, D. Kemp. *Internet X.509 Certificate Request Message Format*. RFC 2511, Network Working Group, The Internet Society, March 1999
19. R. Housley. *Internet X.509 PKI Representation of Key Exchange Algorithm (KEA) Keys in Certificates*. RFC 2528, Network Working Group, The Internet Society, March 1999
20. S.Boeyen, T. Howes, P. Richard. *Internet X.509 PKI Operational Protocols – LDAPv2*. RFC 2559, Network Working Group, The Internet Society, April 1999
21. D. Bryson. *Keeping Secure and Accurate Time on your Network*. Quality Software Solutions, 1998
22. W. LeFebvre. *Daemons and Dragons: Keeping Time with NTP*. UNIX Review, Miller Freeman Inc, March 1998
23. P. H. Dana, B. Penrod. *The Role of GPS in Precise Time and Frequency Dissemination*. Datum Inc, 1998-1999
24. TrueTime Inc. *Network Time Service*. TrueTime Inc, 1998
25. Curry. *Version 3 X.509 Certificates*. Entrust Technologies White Paper, July 1996
26. T. Patterson. *Establishing Electronic Commerce via Certificate Authorities*. Cyberguard-ecom.se-com.com white paper, 1997
27. D.L. Mills. *Measured Performance of the Network Time Protocol*. RFC 1128, Network Working Group, October 1989
28. J. Crowcroft, J. Onions. *NTP over the OSI Remote Operations Seervice*. RFC 1165, Network Working Group, June 1990
29. www.cryptography.com/des *Record-Breaking DES Key Search Completed*. Cryptography Research Inc, 1998

APPENDICES

APPENDIX A BASELINE RESULTS

Table A-1 DWAN Link Capacities and Propagation Delays

Link	Capacity bps	PD	Link	Capacity bps	PD
Cold Lake - B BAR 2	768000	75	North Bay - B BAR 1	384000	35
Cold Lake - T BAR 1	768000	75	North Bay - T BAR 2	384000	35
Shilo - B BAR 2	384000	90	Toronto - B BAR 1	1536000	12
Shilo - T BAR 1	384000	110	Toronto - T BAR 2	1536000	25
Winnipeg - B BAR 2	1536000	37	Petawawa - B BAR 1	768000	45
Winnipeg - T BAR 1	1536000	47	Petawawa - T BAR 2	768000	25
Wainwright - B BAR 2	256000	95	Greenwood - B BAR 2	1536000	67
Wainwright - T BAR 1	256000	85	Greenwood - T BAR 2	1536000	42
Edmonton - B BAR 2	1536000	67	Gander - B BAR 2	384000	130
Edmonton - T BAR 1	1536000	62	Gander - T BAR 2	384000	85
Suffield - B BAR 2	256000	80	St Johns - B BAR 2	384000	130
Suffield - T BAR 1	256000	90	St Johns - T BAR 2	384000	80
Yellowknife - B BAR 2	56000	171	Goose Bay - B BAR 2	192000	670
Yellowknife - T BAR 1	56000	171	Goose Bay - T BAR 2	192000	620
Moose Jaw - B BAR 2	384000	80	Millcove - B BAR 2	56000	146
Moose Jaw - T BAR 1	384000	80	Millcove - T BAR 2	56000	146
Aldergrove - B BAR 1	256000	75	Shearwater - B BAR 2	1536000	67
Aldergrove - T BAR 3	256000	105	Shearwater - T BAR 2	1536000	52
Cornox - B BAR 1	384000	75	Halifax - B BAR 2	1536000	67
Cornox - T BAR 3	384000	105	Halifax - T BAR 2	1536000	52
Masset - B BAR 1	56000	171	St Jean - B BAR 3	1536000	42
Masset - T BAR 3	56000	181	St Jean - T BAR 1	1536000	37
Esquimalt - B BAR 1	1536000	77	Bagotville - B BAR 3	384000	90
Esquimalt - T BAR 3	1536000	92	Bagotville - T BAR 1	384000	60
Kingston - B BAR 1	1536000	27	Gagetown - B BAR 3	768000	50
Kingston - T BAR 2	1536000	24	Gagetown - T BAR 1	768000	65
London - B BAR 1	384000	35	Valcartier - B BAR 3	1536000	17
London - T BAR 2	384000	40	Valcartier - T BAR 1	1536000	22
Trenton - B BAR 1	384000	25	L. Pointe - B BAR 3	1536000	32
Trenton - T BAR 2	384000	70	L. Pointe - T BAR 1	1536000	19

Table A-2 ISAP CPU Utilization

DWAN Baseline ISAP CPU Utilization					
ISAP	Min	Max	Avg	Std Dev	Confidence
Winnipeg	0.0080701	0.0125225	0.01101982	0.00164724	0.0024239855
Shilo	0.0044524	0.0072352	0.00567686	0.00103526	0.0015234362
Cold Lake	0.0097397	0.0116877	0.01063022	0.00064427	0.0009480755
Yellowknife	0.0002783	0.0016697	0.00105746	0.00047877	0.0007045368
Aldergrove	0.0030611	0.0052873	0.00450812	0.00081416	0.0011980693
Masset	0.0061221	0.0128008	0.0094058	0.00214401	0.0031550153
Kingston	0.0097397	0.0116877	0.01057456	0.00065854	0.0009690663
Trenton	0.0019479	0.0061221	0.00422982	0.00158790	0.0023366697
Toronto	0.0030611	0.0055656	0.00422984	0.00085136	0.0012528216
Greenwood	0.0041742	0.0055656	0.00467506	0.00047878	0.0007045539
St Johns	0	0.0002783	0.00022264	0.00011132	0.0001638126
Millcove	0.0002783	0.0013914	0.00094614	0.00037747	0.0005554635
Wainwright	0.0036176	0.0052873	0.00456376	0.00077919	0.0011466168
Edmonton	0.0041742	0.0061221	0.00506466	0.00064425	0.0009480450
Suffield	0.0019479	0.0047307	0.00300538	0.00103226	0.0015190198
Moose Jaw	0.0058438	0.010018	0.00823702	0.00139251	0.0020491469
Comox	0.0083483	0.0116877	0.01046324	0.00120146	0.0017679996
Esquimalt	0.0033393	0.0072352	0.0056212	0.00147673	0.0021730734
London	0.0041742	0.0066787	0.00545426	0.00081796	0.0012036678
North Bay	0.0005566	0.0016697	0.00116878	0.00044525	0.0006552100
Petawawa	0.0038959	0.0072352	0.00523162	0.00113242	0.0016664079
Gander	0	0.0008348	0.00044524	0.00033391	0.0004913692
Goose Bay	0.0033393	0.0052873	0.00422982	0.00071275	0.0010488481
Shearwater	0.0038959	0.0058438	0.00506464	0.00077518	0.0011407193
Halifax	0.0033393	0.0069569	0.00517596	0.00141457	0.0020816096
Bagerville	0.0025045	0.005009	0.0035063	0.00083668	0.0012312180
Valcartier	0.0027828	0.0061221	0.00434116	0.00132408	0.0019484433
St Jean	0.0008348	0.0019479	0.00139138	0.00039355	0.0005791234
Gagetown	0.0030611	0.0072352	0.00528728	0.00164157	0.0024156544
Longue Pointe	0.0033393	0.0055656	0.00461942	0.00085502	0.0012582091
Borden MAN	0.0005566	0.0013914	0.0010018	0.00028378	0.0004175957
NDHQ MAN	0.0186991	0.0228268	0.02038858	0.00173918	0.0025592798
Inf	0.0361946	0.0501065	0.04604554	0.00501063	0.0073733736

Table A-3 DWAN Baseline Results

DWAN Link Utilization					
Link # Traffic Direction	Min	Max	Avg	Std Dev	Confidence
9					
FROM B BAR 2	2.1859	2.2166	2.20064	0.01121171	0.0164985553
FROM Cold Lake	1.8005	1.8275	1.8154	0.01044911	0.0153763633
10					
FROM T BAR 1	12.63	12.7	12.662	0.02561250	0.0376899924
FROM Cold Lake	4.7559	4.7958	4.77514	0.01465082	0.0215593635
5					
FROM B BAR 2	2.9089	2.9557	2.93206	0.01839506	0.0270692016
FROM Shilo	3.8303	3.8774	3.84828	0.01601404	0.0235654178
6					
FROM T BAR 1	29.75	30.12	29.936	0.12595237	0.1853448321
FROM Shilo	15.48	15.6	15.546	0.04800000	0.0706342549
1					
FROM B BAR 2	2.6528	2.6781	2.66216	0.00920339	0.0135432217
FROM Winnipeg	2.5706	2.609	2.59282	0.01354642	0.0199341991
2					
FROM T BAR 1	10.48	10.52	10.5	0.01264911	0.0186137605
FROM Winnipeg	8.7274	8.7753	8.74586	0.01716445	0.0252582969
3					
FROM B BAR 2	2.6211	2.6923	2.66104	0.02455831	0.0361387011
FROM Wainwright	1.2563	1.2841	1.26716	0.01040492	0.0153113314
4					
FROM T BAR 1	42.1	42.66	42.444	0.18596774	0.2736602642
FROM Wainwright	18.44	18.62	18.528	0.06431174	0.0946377477
7					
FROM B BAR 2	3.7253	3.7387	3.73118	0.00495072	0.0072852128
FROM Edmonton	1.3264	1.3393	1.33302	0.00490689	0.0072207237
8					
FROM T BAR 1	12.98	13.03	13.012	0.01720465	0.0253174515
FROM Edmonton	3.7555	3.827	3.79608	0.02367154	0.0348337767
11					
FROM B BAR 2	1.8641	1.918	1.90322	0.01976840	0.0290901278
FROM Suffield	1.4018	1.456	1.43448	0.01801371	0.0265080144
12					
FROM T BAR 1	31.55	31.79	31.67	0.08508819	0.1252112681
FROM Suffield	23.92	24.19	24.08	0.09674709	0.1423678922

Table A-3, Continued

13					
FROM B BAR 2	11.92	12.3	12.144	0.13573504	0.1997404818
FROM Yellowknife	3.6299	3.7661	3.71292	0.05151592	0.0758080904
14					
FROM T BAR 1	25.89	26.64	26.334	0.26657832	0.3922825189
FROM Yellowknife	12.01	12.22	12.116	0.08546344	0.1257634704
15					
FROM B BAR 2	3.173	3.2633	3.21632	0.03284201	0.0483285600
FROM Moose Jaw	2.443	2.5032	2.47464	0.02182050	0.0321098969
16					
FROM T BAR 1	12.82	12.98	12.888	0.05192302	0.0764071630
FROM Moose Jaw	5.7764	5.8291	5.81032	0.01797425	0.0264499507
17					
FROM Aldergrove	12.31	12.41	12.348	0.03709447	0.0545862610
FROM B BAR 1	5.1589	5.2561	5.20522	0.03560261	0.0523909082
18					
FROM Aldergrove	17.88	17.99	17.948	0.04166533	0.0613124953
FROM T BAR 3	8.2661	8.4735	8.37934	0.08181280	0.1203913799
19					
FROM Comox	2.6887	2.7663	2.71646	0.02726768	0.0401256728
FROM B BAR 1	11.47	11.6	11.516	0.04841487	0.0712447611
20					
FROM Comox	6.4771	6.5914	6.54766	0.03783229	0.0556719965
FROM T BAR 3	15.93	16.26	16.098	0.10684568	0.1572284419
21					
FROM Masset	1.7309	1.8241	1.77986	0.03140997	0.0462212507
FROM B BAR 1	2.3591	2.4821	2.41786	0.05324399	0.0783510315
22					
FROM Masset	6.6423	7.0592	6.85352	0.14388114	0.2117278522
FROM T BAR 3	14.67	15.05	14.858	0.13166624	0.1937530643
23					
FROM Esquimalt	2.2771	2.313	2.29334	0.01170019	0.0172173763
FROM B BAR 1	3.3688	3.4018	3.38406	0.01121920	0.0165095766
24					
FROM Esquimalt	4.9459	4.9641	4.9584	0.00640250	0.0094215788
FROM T BAR 3	13.38	13.43	13.39	0.02000000	0.0294309395
25					
FROM Kingston	3.6881	3.738	3.71684	0.01718669	0.0252910254
FROM B BAR 1	2.8252	2.8697	2.84844	0.01531151	0.0225316092

Table A-3, Continued

26					
FROM Kingston	4.2053	4.2597	4.22502	0.01864504	0.0274370555
FROM T BAR 2	6.9617	7.0282	6.98602	0.02356518	0.0346772660
27					
FROM London	2.1253	2.1873	2.15956	0.02089408	0.0307466139
FROM B BAR 1	5.3212	5.3709	5.34156	0.02028138	0.0298450043
28					
FROM London	6.3302	6.4199	6.38116	0.03441021	0.0506362372
FROM T BAR 2	22.87	22.98	22.924	0.04127953	0.0607447747
29					
FROM Trenton	3.2837	3.38	3.34196	0.03355465	0.0493772408
FROM B BAR 1	20.3	20.49	20.416	0.07337575	0.1079758562
30					
FROM Trenton	14.34	14.51	14.404	0.05678028	0.0835548468
FROM T BAR 2	33.62	33.77	33.684	0.05817216	0.0856030636
31					
FROM North Bay	1.544	1.5856	1.56006	0.01487247	0.0218855400
FROM B BAR 1	2.263	2.3203	2.29096	0.02018718	0.0297063897
32					
FROM North Bay	18.09	18.29	18.182	0.06554388	0.0964508961
FROM T BAR 2	19.04	19.2	19.134	0.05276362	0.0776441511
33					
FROM Toronto	1.098	1.1089	1.10456	0.00361972	0.0053265935
FROM B BAR 1	1.86	1.8887	1.86992	0.01014306	0.0149259845
34					
FROM Toronto	2.0015	2.0196	2.00912	0.00639606	0.0094121046
FROM T BAR 2	4.924	4.966	4.94188	0.01576266	0.0231955017
35					
FROM Petawawa	5.1684	5.211	5.18956	0.01406309	0.0206944923
FROM B BAR 1	11.5	11.64	11.57	0.04472136	0.0658095814
36					
FROM Petawawa	7.0263	7.1089	7.0629	0.03214312	0.0473001073
FROM T BAR 2	19.84	19.91	19.87	0.02449490	0.0360453922
37					
FROM Greenwood	0.9419	0.9523	0.9455	0.00365677	0.0053811155
FROM B BAR 2	1.4353	1.4662	1.44982	0.01189914	0.0175101476
38					
FROM Greenwood	2.2774	2.311	2.29426	0.01393293	0.0205029537
FROM T BAR 2	6.135	6.2127	6.17394	0.02819628	0.0414921556

Table A-3, Continued

39					
FROM Gander	0.9621	1.0035	0.98524	0.01356549	0.0199622493
FROM B BAR 2	1.3226	1.3608	1.33824	0.01493782	0.0219817006
40					
FROM Gander	2.7109	2.7771	2.74982	0.02286993	0.0336541732
FROM T BAR 2	5.7914	5.8848	5.84118	0.04024875	0.0592279213
41					
FROM St Johns	0.6423	0.6753	0.6565	0.01140666	0.0167854430
FROM B BAR 2	2.305	2.3394	2.32586	0.01134506	0.0166947896
42					
FROM St Johns	4.6756	4.7398	4.71356	0.02429062	0.0357447939
FROM T BAR 2	14.24	14.32	14.27	0.02828427	0.0416216338
43					
FROM Goose Bay	1.5763	1.6632	1.6219	0.03194489	0.0470084064
FROM B BAR 2	2.349	2.4282	2.38506	0.02582933	0.0380090756
44					
FROM Goose Bay	6.5709	6.6599	6.62782	0.03334699	0.0490716592
FROM T BAR 2	10.77	10.85	10.81	0.02607681	0.0383732504
45					
FROM Millcove	2.0563	2.2689	2.1828	0.07251353	0.1067070655
FROM B BAR 2	3.121	3.3036	3.22618	0.05895800	0.0867594646
46					
FROM Millcove	6.8419	7.0737	6.94636	0.07521924	0.1106886485
FROM T BAR 2	10.06	10.53	10.302	0.16166632	0.2378995888
47					
FROM Shearwater	1.7395	1.7695	1.75528	0.01034300	0.0152202069
FROM B BAR 2	1.4939	1.5139	1.50332	0.00663578	0.0097648646
48					
FROM Shearwater	2.9077	2.9412	2.91836	0.01201043	0.0176739102
FROM T BAR 2	4.6659	4.7149	4.69424	0.01588579	0.0233766898
49					
FROM Halifax	1.3183	1.344	1.32836	0.00845898	0.0124477914
FROM B BAR 2	5.1907	5.2606	5.23274	0.02385167	0.0350988602
50					
FROM Halifax	8.7242	8.7859	8.7561	0.02203198	0.0324210888
FROM T BAR 2	15.88	15.96	15.912	0.02712932	0.0399220687
51					
FROM St Jean	4.0574	4.1076	4.08388	0.01604873	0.0236164539
FROM B BAR 3	2.4504	2.4685	2.45818	0.00824437	0.0121319745

Table A-3, Continued

52	FROM St Jean	0.03	0.0328	0.031	0.00095708	0.0014083865
	FROM T BAR 1	5.0683	5.0976	5.08484	0.01084815	0.0159635652
53						
	FROM Bagoiville	4.1866	4.2499	4.21678	0.02474473	0.0364130307
	FROM B BAR 3	4.1392	4.2584	4.193	0.04318944	0.0635552945
54						
	FROM Bagoiville	8.1087	8.1783	8.14282	0.02503161	0.0368351930
	FROM T BAR 1	19.1	19.23	19.162	0.04261455	0.0627093144
55						
	FROM Gagetown	8.3781	8.4727	8.42872	0.03215167	0.0473126863
	FROM B BAR 3	4.0515	4.0642	4.05858	0.00432823	0.0063691980
56						
	FROM Gagetown	17.58	17.74	17.676	0.05748043	0.0845851551
	FROM T BAR 1	25.92	26.02	25.954	0.03555278	0.0523175825
57						
	FROM Valcartier	4.1338	4.1761	4.15444	0.01452551	0.0213749683
	FROM B BAR 3	7.0873	7.1534	7.12658	0.02167684	0.0318984954
58						
	FROM Valcartier	6.3585	6.4123	6.38848	0.01854102	0.0272839890
	FROM T BAR 1	12.03	12.13	12.072	0.03487119	0.0513145965
59						
	FROM Longue Pointe	3.8096	3.8209	3.81372	0.00403158	0.0059326526
	FROM B BAR 3	4.1811	4.2279	4.19686	0.01703955	0.0250745041
60						
	FROM Longue Pointe	5.7064	5.7593	5.73236	0.01701230	0.0250344007
	FROM T BAR 1	12.45	12.49	12.466	0.01624808	0.0239098083
	FROM B BAR 3 (66)	50.64	50.74	50.698	0.03762978	0.0553739823
	FROM T BAR 3 (66)	28.88	28.98	28.94	0.04195235	0.0617348596
	FROM Peakes BAR	31.57	31.64	31.614	0.03200000	0.0470895032
	FROM B BAR 1	43.93	44.02	43.98	0.03521363	0.0518185162

APPENDIX B SIMULATION RESULTS

B.1 Simulation Results Introduction

This appendix contains all of the results from the PKI simulations in Excel graph format. Section B.2 covers the Option 1 results, B.3 the Option 2 and 3 results, and B.4 presents the Variation 1 and 2 Excel bar graphs.

B.2 Option 1 Results

Figure B. 1 Option 1 Wainwright, Winnipeg, Shilo, Edmonton - Borden

Figure B. 2 Option 1 Cold Lake, Yellowknife, Moose Jaw, Suffield - Borden

Figure B. 3 Option 1 Aldergrove, Esquimalt, Comox, Masset - Borden

Figure B. 4 Option 1 Kingston, Trenton, London, North Bay - Borden

Figure B. 5 Option 1 Toronto, Petawawa - Borden

Figure B. 6 Option 1 Wainwright, Winnipeg, Shilo, Edmonton - Tunney's

Figure B. 7 Option 1 Cold Lake, Yellowknife, Suffield, Moose Jaw - Tunney's

Figure B. 8 Option 1 Aldergrove, Masset, Comox, Esquimalt - Tunney's

Figure B. 9 Option 1 Kingston, Trenton, London, North Bay - Tunney's

Figure B. 10 Option 1 Toronto, Petawawa - Tunney's

Figure B. 11 Option 1 Greenwood, St John's, Gander, Goose Bay - Tunney's

Figure B. 12 Option 1 Millcove, Shearwater, Halifax - Tunney's

Figure B. 13 Option 1 Bagotville, Valcartier, Gagetown, Longue Pointe - Tunney's

B.3 Option 2 & 3 Simulation Results

Since Option 2 and 3 have nearly identical results, only the Option 3 graphs are presented in this section.

Figure B. 14 Option 3 Wainwright, Winnipeg, Shilo, Edmonton

Figure B. 15 Option 3 Cold Lake, Suffield, Yellowknife, Moose Jaw

Figure B. 16 Option 3 Aldergrove, Comox, Masset, Esquimalt

Figure B. 17 Option 3 Kingston, London, Trenton, North Bay

Figure B. 18 Option 3 Toronto, Petawawa

Figure B. 19 Option 3 Greenwood, Gander, St. John's, Goose Bay

Figure B. 20 Option 3 Mill Cove, Shearwater, Halifax, St. Jean

Figure B. 21 Option 3 Bagotville, Gagetown, Valcartier, Longue Pointe

B.4 Variation 4 and 5 Simulation Results

The results from the Variation 1 and 2 simulations are presented together in this section. In the following bar graphs, the columns labeled 'B' represent the baseline model link utilization. The bars labeled 1, 2, and 3 are the link utilization for Option 1, 2, and 3 respectively, extracted from the models using the 5 times traffic multiplier. The bars labeled 4 and 5 are the link utilization for Variation 1 and 2 respectively, both extracted from the Variation simulation results. Not every base is represented in the bar charts. A set of western Canadian bases, one of a mix of eastern and western bases, and one of just eastern bases are presented.

Figure B. 22 Variation 1 and 2 - Winnipeg, Cold Lake, Edmonton, Moose Jaw - Borden

Figure B. 23 Variation 1 and 2 - Esquimalt, Kingston, Trenton, Toronto - Borden

Figure B. 24 Variation 1 and 2 - Greenwood, Shearwater, Halifax, Gagetown - Tunney's

VITA

Vita

Name: David Edward DePlanché

Place and year of birth: Niagara Falls, Ontario, 1960

Education: Royal Military College of Canada, 1998-2000;
University of Manitoba, 1990-1994.
Awarded Bachelor of Computer Science (Honours) May 1994.

Experience: Post Graduate Student, Royal Military College of Canada,
1998-2000

Maintenance Planning and Production Officer, 402 Squadron,
Winnipeg, 1995-1998

Aerospace Engineering Officer Basic Course, Canadian Forces
School of Aerospace and Technical Engineering, 1994-1995

Radar and Communications Systems Technician, Canadian
Forces, 1980-1989