Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Long-term operating system maintenance

A Linux case study

R. Carbone
Certified Hacking Forensic Investigator (EC-Council CHFI)
Certified Incident Handler (SANS)
DRDC Valcartier

Defence R&D Canada – Valcartier

Technical Memorandum
DRDC Valcartier TM 2007-150
March 2013

Canada

Long-term operating system maintenance

A Linux case study

R. Carbone
Certified Hacking Forensic Investigator (EC-Council CHFI)
Certified Incident Handler (SANS)
DRDC Valcartier

Defence R&D Canada – Valcartier

Technical Memorandum
DRDC Valcartier TM 2007-150
March 2013

Principal Author

Original signed by Richard Carbone

Richard Carbone

Programmer/Analyst

Approved by

Original signed by Guy Turcotte

Guy Turcotte

Head/Mission Critical Cyber Security Section

Approved for release by

Original signed by Christian Carrier

Christian Carrier

Chief Scientist

Abstract

In TM 20 06-595, *Operating system hardware re configuration: A case study for Linux* , it was determined t hrough exper imentation that a Li nux-based C2 operating sy stem can successf ully undergo a hardware migration and operating system hardware reconfiguration. The direct benefit of this is t he ability to for go any new operating system reinstallation i n or der to support ne wer hardware by using mechanisms internal to the operating system that support changes in hardware. This results in a decreas ed waiting time for sy stem r eaccreditation and redeploy ment. Since a n operating sy stem can evol ve over ti me, it can accommodate cha nges in the s ystem's hardware, thus presenting a tangible advantage for the Roy al Navy, as this allows the operating system to be maintained over the lon g-term. Howev er, there are co mplexities involved when m aintaining an operating sy stem for long periods. Therefore, th is report serves as an introduction and a sim ple methodology for perform ing s ystem maintenance-related tasks that include up grading, up dating, as well as data backups and restoration. This repor t is neither all-inclusive nor a replacement for qualified s ystem ad ministrators with years of expe rience. Instead, it can be used as a source of information to pro vide recommended practices, pro cedures, and i nformation for helpi ng to plan for long-term system maintenance.

Résumé

Dans le TM 2006- 595, *Operating system hardware reconfiguration: A case study for Linux* , il a été déterminé expérimentalement qu'un système d'exploitation d'un C2 basé sur Linux peut subir avec succ ès une m igration m atérielle ainsi qu'une reconfiguration m atérielle du sy stème d'exploitation. Le bénéfice direct est la capac ité de ne pas avoir à procéd er à une nou velle réinstallation du système d'exploitation afin de supporter du matériel plus récent et ce, en utilisant les mécanismes internes du sy stème d'exploitation qui soutiennent les changements du matériel. Il en résulte une dim inution du tem ps d'attente pour la ré-accréditation et le redéploiem ent du système. Étant donné qu 'un sy stème d'exploitation peut évoluer au fil du tem ps, il peut donc s'adapter aux changements dans le matériel du système, ce qui présente un avant age tangible pour la Marine roy ale canadienne. Cela p ermet au sy stème d 'exploitation d 'être maintenu à long terme. Cependant, maintenir un sy stème d'e xploitation sur une longue période engendre des complexités. Par conséquent, le présent ra pport se veut une int roduction et une m éthodologie simple pour effectuer les tâches reliées à la maintenance d'u n système qui incluent la m ise à niveau, la mise à jour, ainsi que la sauvegarde de données et leur restauration. Ce rapport n' est ni exhaustif, ni un rem placement pour les ad ministrateurs de sy stèmes qualifiés avec plusieurs années d'expérience. Il doit plutôt être utilisé comme une source d'informations utile pour fournir des pratiques reco mmandées, des procédures, ai nsi que des inform ations pour aider à la planification de la maintenance à long terme d'un système.

This page intentionally left blank.

Executive summary

Long-term operating system maintenance: A Linux case study

Carbone, R.; DRDC Valcartier TM 2007-150; Defence R&D Canada – Valcartier; March 2013.

The Canadian Royal Navy's Directorate of Maritime Ship Support (DMSS), under the auspice of the Halifax Modernized Command Control System (HMCCS) project, requested that DRDC Valcartier performs an evaluation verifying if the Linux operating system has the ability to withstand various hardware upgrades over its expected lifetime as the new Halifax-class C2 operating system. This has been examined in *Operating system hardware reconfiguration: A case study for Linux*, which looked at Linux-based hardware migration and operating system hardware reconfiguration. Prior to this study, other work mandated by the Royal Navy requested that DRDC Valcartier examines various support strategies appropriate for the long-term maintenance of a free and open source software-based operating system such as Linux. These findings were laid out in *Life-cycle support for information systems based on free and open source software*.

Once the long-term support strategies and necessary reconfiguration capabilities have been established, an appropriate methodology could be developed for upgrading and updating an operating system. This would ensure a smooth transition to newer hardware via operating system hardware reconfiguration or hardware migration. Therefore, this study proposes a methodology for upgrading and updating an operating system so that installing a newer operating system is not required in order to support newer hardware. Developing a suitable methodology for maintaining an operating system so that it can adapt to periodic hardware changes was not trivial. Many factors had to be examined and taken into account. It is highly suggested that the reader have some basic knowledge of Linux, UNIX and system administration before reading this study.

This study should not be construed as authoritative, but rather as a guide for performing the necessary system administration related tasks required for system updating and upgrading. Additionally, a section on backup-related issues has been included, so that readers can familiarize themselves with it prior to undertaking any system maintenance-related task.

Sommaire

Long-term operating system maintenance: A Linux case study

Carbone, R. ; DRDC Valcartier TM 2007-150 ; R & D pour la défense Canada – Valcartier; mars 2013.

Le Directeur - Soutien aux navires (DSN) de la Marine royale canadienne, sous l'égide du projet de modernisation du système de commandement et contrôle de la classe Halifax, a demandé que RDDC Valcartier effectue une évaluation pour vérifier si le système d'exploitation Linux a la capacité de supporter plusieurs mises à niveau matérielles au cours de sa durée de vie prévue comme le nouveau système d'exploitation de C2 de la classe Halifax. Ceci a été examiné dans *Operating system hardware reconfiguration: A case study for Linux*, qui s'est penché sur la migration matérielle basée sur Linux et la reconfiguration matérielle de système d'exploitation. Avant cette étude, d'autres travaux mandatés par la Marine royale ont demandé que RDDC Valcartier examine diverses stratégies de soutien appropriées pour le maintien à long terme d'un système d'exploitation basé sur des logiciels libres comme Linux. Ces conclusions ont été présentées dans *Life-cycle support for information systems based on free and open source software*.

Une fois que les stratégies de soutien à long terme et les capacités de reconfiguration nécessaires ont été mises en place, une méthodologie appropriée pourrait être développée afin de mettre à jour et à niveau un système d'exploitation. Cela permettrait d'assurer une transition en douceur vers du matériel plus récent par le biais de la reconfiguration matérielle du système d'exploitation ou la migration matérielle. Par conséquent, cette étude propose une méthodologie pour la mise à niveau et à jour d'un système d'exploitation afin que l'installation d'un nouveau système d'exploitation ne soit pas nécessaire pour soutenir du matériel plus récent. Développer une méthodologie appropriée pour maintenir un système d'exploitation afin qu'il puisse s'adapter à des changements périodiques de matériel n'a pas été anodin. De nombreux facteurs ont dû être examinés et pris en considération. Il est fortement suggéré que le lecteur ait une connaissance de base de Linux, d'UNIX et de l'administration de systèmes avant de lire cette étude.

Cette étude ne doit pas être interprétée comme faisant autorité, mais plutôt comme un guide pour effectuer les tâches nécessaires liées à l'administration de systèmes qui sont requises pour la mise à jour et à niveau de ceux-ci. De plus, une section sur les questions liées à la sauvegarde de données a été incluse, afin que les lecteurs puissent se familiariser avec celles-ci lui avant d'entreprendre toute tâche liée à la maintenance de systèmes.

This page intentionally left blank.

Table of contents

Acknowledgements

The author would like to extend a large thanks to Philippe Charland who helped provide the necessary comments and corrections in order to reincarnate this document as a Technical Memorandum. Moreover, the author would also thank him for peer reviewing this document.

1 Introduction

1.1 Objective

The objective of this technical memorandum is to examine an update and upgrade-based methodology that can be applied to the Linux operating system in order to maintain it and its hardware throughout its expected long-term service. It is expected that over the years, the Royal Navy's frigate C2 systems will not only require operating system maintenance, but that they will also periodically undergo hardware upgrades. This technical memorandum addresses the necessary methodologies required to proceed with an operating system update or upgrade and thus enable the system to support and accommodate periodic hardware changes. This is accomplished through two methodologies. The first provides an outline for maintaining system integrity via backups and restorations, while the other examines the various issues that must be assessed prior to performing any system-related maintenance. Both methodologies, when combined together, form a comprehensive system maintenance methodology that will help the reader succeed in maintaining his system over the long-term.

1.2 Background

1.2.1 Reports

This technical memorandum is the fourth and final report in a series of reports prepared for the Royal Navy. The very first report [2] examined the various long-term support strategies needed for maintaining an open source operating system such as Linux. Specifically, it was determined that long-term support could be achieved by (in order of preference): (1) the maintainer of the distribution (e.g. vendor), (2) a sub-contractor, (3) the Royal Navy and (4) a consortium.

The second report [3] was an in-depth analysis of two licenses that compared Red Hat and Novell Suse's enterprise Linux operating systems. They were examined because the Royal Navy had expressed interest in these two operating systems due to: (1) both companies are North American and (2) are already involved with the military establishment, particularly in the U.S.

The third report [1] examined, through experimentation and observation, the Linux operating system's ability to adapt to changes in its underlying hardware. It was found that Linux can reconfigure itself in order to take advantage of newer hardware, but only if the kernel is recent enough to support that newer hardware. It was concluded that both long-term system maintenance and hardware adaptation could be accomplished in so long as the operating system's kernel is kept up to date.

Reports [5, 6 and 7] are informal internal reports that were the precursor to both this technical memorandum and reports [1, 3]. These internal reports were not rigorous enough to be released as formal reports from DRDC.

This technical memorandum is the final report in this series. It examines the methodology necessary for maintaining an operating system so that it can be supported over the long-term with the added benefit of successfully accommodating periodic hardware upgrades.

1.2.2 Mandates

All the reports mandated, including this one, have been conducted for the Royal Navy's Directorate of Maritime Ship Support (DMSS 8), under the auspice of the Halifax Modernized Command Control System (HMCCS). These reports have been useful for providing technical information to the Royal Navy, aiding it to determine which new operating system, if any, would replace the current C2 operating system aboard the Canadian Halifax-class frigate. These frigates, including both the computer operating system and hardware, are currently undergoing modernization.

The original mandate of DRDC Valcartier was to examine and address the various long-term support strategies available to the Royal Navy should they decide to pursue the deployment of a Linux-based operating system aboard the frigates. Furthermore, the Royal Navy wished to understand through which means long-term operating system maintenance could be achieved and supported. This work was carried out in [2].

Report [3] was a direct response to a question posed by the HMCCS project that was to conduct a comparison between the corporate licenses of two popular North American Linux vendors. The conclusion that was reached may have an important impact on determining the outcome of which vendor's distribution could potentially be used as the new C2 operating system.

However, as work was underway for reports [2, 3], the HMCCS project requested that DRDC Valcartier examines whether the Linux operating system could adapt to periodic changes in hardware and determine how this would affect overall long-term support. This was studied in Report [1] and consisted of a series of short-lived experiments and observations that was the first of its kind as none could be found in the public literature. Since the new C2 operating system could be in use for 15 to 25 years, the Royal Navy needed to ascertain the maintainability of the operating system and its ability to adapt to periodic hardware changes that are bound to occur.

This technical memorandum was the result of previous reports that examined the Royal Navy's potential use and implementation of Linux. Specifically, this technical memorandum uses the results obtained in Report [1] and goes one step further by examining how the Linux operating system can be maintained over the long-term using a vendor-neutral technical methodology. The outcome is a report that provides a global overview of Linux system maintenance that can be understood by IT managers, system administrators and others technical support staff.

1.3 Particulars

1.3.1 Reader

It is assumed that the reader is familiar with Linux, UNIX or BSD-based operating systems and has system administration-related experience. For reasons of brevity, it is not possible to provide

all the required background, contextual and technical information that a full in-depth examination would normally provide.

This text is vendor neutral and does not directly examine the implementation of various low-level commands required to carry out actions examined throughout the methodology. Finally, it is also assumed that the reader will be using a commercially supported Linux-based operating system.

1.3.2 Royal Navy

The Royal Navy prefers that changes to the operating system be as minimal as possible in order to accelerate the process of reaccreditation and recertification. However, where operating system maintenance is concerned, this is not always easy. Much will depend on what requires maintenance and how it will be performed. There are many variables to consider prior to performing any system maintenance-related action. Guided by the two methodologies, the reader can perform the appropriate set of actions necessary to maintain the system and support the newer hardware.

However, the Royal Navy must understand that operating system maintenance is not always simple. The main method of system maintenance employs operating system updates, upgrades and manual system maintenance. Updates and upgrades are generally provided by the distribution's vendor, while manual system maintenance is performed by the customer or support provider. Each method has its own advantages and disadvantages. Ultimately, the maintenance option that will be used at any point in time will depend on what requires maintenance as well as who is providing maintenance support.

1.3.3 Report

The technical memorandum is broken down into various main topics or sections. The first is the introduction that presents an overview to the reader. The second section examines technical information about operating systems and hardware support. The third section is a broad overview for performing backup and restoration-related tasks. While this section goes in-depth, it is not meant as a replacement for various books already written on the subject; instead, it presents material that should be considered prior to undertaking any specific task or action. In addition, this section does present some new information not found in existing books on the topic. Finally, the fourth section provides a broad overview of operating system maintenance that examines the higher-level issues surrounding updates and upgrades. Manual system maintenance has been deliberately excluded from this report, as it is too technically complex and cumbersome to examine in a short report. Furthermore, it varies considerably from software package to software package and its success is highly dependent on the experience and knowledge of the support staff.

This technical memorandum is not meant to be an all-encompassing authority on the various subjects of interest. Instead, it should be considered a guide for gaining a higher-level overview on how to approach the subject matter from a global perspective. Certain topics that could have been explored more in-depth were left out to remain vendor neutral.

Finally, in order to ascertain this study's relevancy, it is important to understand the correlation between the operating system and its life expectancy. As a function of time, hardware is expected

to change. It is therefore only reasonable to expect that a C2 system will undergo hardware changes. In order for the operating system to support newer hardware, it must be kept up-to-date. This is particularly important for hardware employing technology that is not currently supported by older kernels and can be accomplished through updates, upgrades and manual system maintenance.

1.3.4 Methodology

A generic methodology will help to accommodate for the various types of long-term system maintenance and allow the reader to use whichever form is most suitable to his current requirements. However, due to the generalities involved in formulating a generic methodology, many specifics are kept out of the overall process so that the reader can decide which tools and other specifics are most appropriate. Furthermore, a generic approach keeps the text short and concise. In addition, in order to maintain brevity and simplicity, some issues, concepts and tasks have been altogether left out. They should not affect the overall clarity, as they are minor details that can be filled in by the reader's experience and knowledge.

It is important to understand the difference between an operating system update and upgrade. They only differ in the magnitude of changes made to the operating system. An update generally has less overall impact on the operating system than an upgrade. However, in order to ensure a minimal impact for either, they should be performed periodically.

Furthermore, a generic methodology is the best possible approach in order to ensure a successful hardware migration and operating system hardware reconfiguration whenever newer hardware must be supported. As examined in [1], a migration and reconfiguration are generally the most appropriate methods for either adapting to newer hardware or moving from one platform to the next. By avoiding the necessity of installing a newer operating system, many potential issues of conflict and contention can be avoided. In addition, employing a generic methodology frees the reader to focus on higher-level issues.

A generic methodology allows for a global overview of the concepts and tasks to be accomplished. How they are completed and through which means is of little concern here. Finally, a generic approach facilitates the interchanging of various lower-level technologies and tools.

2 Technical background

2.1 Objective

In this section, several important aspects concerning operating systems, regardless of the underlying system type or computer platform, are examined. The first aspect presents basic technical information on operating system migrations and reconfigurations. The second presents technical information on operating system dependencies and call facilities. Finally, compatibility-based issues are briefly examined.

2.2 Reconfigurations and migrations

2.2.1 Background

The most common reason for performing a hardware reconfiguration is to save time and avoid a complete security recertification of the operating system. Rather than install a newer operating system, often considered a daunting task, it may be more appropriate to leverage on the existing operating system with its current configurations and settings.

The Royal Navy has stated that it may not be possible for them to perform full operating system upgrades[1] due to restraints in security, certification, auditing, system and change management, as well as overall maintenance. Essentially, the Royal Navy is opting to freeze or "lock out" the operating system with a given configuration for its entire lifetime of between 15 and 25 years. However, it is reasonable to assume that over this period, there will be periodic hardware upgrades to the C2 systems. These upgrades may include minor hardware changes or complete system overhauls.

Generally, outdated operating systems that are not kept at least partially up to date will not fare well in supporting modern hardware. Although the Royal Navy may not be able to switch to a more recent operating system, it can nevertheless benefit from a reconfiguration in so long as the kernel, drivers and other subsystems are kept reasonably up to date. Thus, existing applications, preferences and configurations can be maintained, as can the majority of operating system dependencies and interdependencies, thereby forgoing the necessity for an in-depth recertification of the system before redeployment. Only software that has actually changed would need to be tested and recertified, specifically the kernel. In large modern operating systems such as Linux, the kernel normally represents less than 1% to 2%[2] of the entire operating system. Furthermore, using modern static analysis techniques it is possible to better assess various vulnerabilities and flaws in the kernel source code [8, 9]. This will help to make a multi-month recertification process last only several weeks, significantly reducing costs and complexities.

[1] By this it is meant going from one version of a Linux distribution to another version of that distribution. An example of this would be going from Red Hat Enterprise Linux (RHEL) 3 to version RHEL 4.

[2] This was the case when this document was originally written in 2007. It continues to remain the case in 2013.

2.2.2 Reconfigurations

A hardware reconfiguration, with respect to reports [1, 5, 6 and 7], is also commonly known as *operating system hardware reconfiguration*, *operating system reconfiguration* or *reconfiguration*, although they all have the same meaning. The author defines them as "*the ability for an operating system to effectively deal with any changes to the underlying hardware and effectively perpetuate those changes to the appropriate software layers of the operating system, such that the changes should remain as transparent as possible to the user.*"

Several things must be present for a successful reconfiguration prior to changing any hardware. Firstly, the operating system must actually provide software-based support for hardware detection/redetection. Subsequently, a mechanism must exist by which hardware detection/redetection can be triggered, either automatically at system start-up, or manually by the system administrator at any given time. Thirdly, the changes to the operating system must be effectuated in such a manner that they are transparent to the end-user and do not generally require the administrator to manually modify system configuration files.

The majority of today's Linux-based distribution kernels are built "out of the box" to support a wide variety of hardware devices and computing platforms. Most of these operating systems also support various mechanisms for detecting hardware during their initial setup and installation phases. They also have redetection-based tools to detect post-installation changes to the hardware and make the necessary operating system changes to various configuration files. Some operating systems detect hardware changes automatically, sometimes referred to as *dynamic reconfiguration*, while others do not have this ability where redetection must be manually instantiated referred to as *static reconfiguration*. So long as the operating system supports either dynamic or static reconfiguration, it can be said that the operating system is capable of hardware reconfiguration. Whether support is dynamic or static is perhaps a reflection of the level of sophistication of the operating system itself. However, the ability for a reconfiguration to occur is commensurate with the maturity of the operating system. Older Linux systems often encountered difficulty supporting existing hardware. In such cases where older Linux systems are used, it is not realistic to expect it to support newer hardware components.

2.2.3 Migrations

A *hardware migration* (with respect to reports [1, 5, 6 and 7]), often referred to as *migration*, is a term similar to operating system hardware reconfiguration. However, they differ in that a migration[3] is an operating system hardware reconfiguration that takes place only after the entire underlying computing platform has been replaced. In other words, the operating system is altogether transferred to another computer by various means (e.g. disk copying, etc.). An example of this would be moving a Linux-based operating system from a Pentium-class system to a Pentium Xeon, i2, i3 or i7-class system. Most, if not all the new system's hardware will be completely different from that of its predecessor. An operating system hardware reconfiguration then takes place only once the power is applied to the new system is allowed to boot. Depending on the maturity of the Linux operating system in question, the reconfiguration, if it occurs may be either dynamic or static in nature.

[3] Migrations can only be performed on systems with the same basic architecture (e.g. x86, x64, etc.).

2.3 Operating systems

2.3.1 Background

Operating systems are far larger and complex than they were, even ten years ago. For this reason, in this section, a brief examination of operating systems and the technical issues, as well as the complexities surrounding them, will be examined in the following subsections. Although this section is specifically concerned with the Linux operating system, it is equally applicable to other open source and proprietary operating systems as well.

2.3.2 Definition

The term operating system has been given many varying definitions over the years. In this text, the more classical definition of the term is used. An operating system can be big or small. It is not defined by size. At a minimum, an operating system is a collection of executable code separated into individual computer files that controls the computer's hardware and user-based applications, tools and utilities. An operating system is comprised of (1) a kernel, (2) a shell or GUI (for interacting with the kernel and launching applications) and (3) user-based applications.

The kernel provides the low-level facility (or layer) that interacts directly with the system's hardware. The kernel consists of various device drivers, memory management components and a general framework for interacting with the hardware and running applications. The shell or GUI accepts user-based input used for interacting with the system's hardware and running various applications and utilities that perform some useful work on the user's behalf. Finally, user applications and utilities exist to provide services and functionality to the user.

With Linux, it is important to understand the difference between the Linux kernel and a Linux-based operating system. Linux, in of itself, refers only to the GNU/Linux kernel which includes drivers and various subsystems such as memory management components. A Linux-based operating system, at its most basic, is no more different from any other operating system. All operating systems require a kernel, a shell (or GUI) for user-based interactions with the kernel and a collection of user-based applications.

2.3.3 System dependencies

2.3.3.1 Interdependencies

Every operating system has interdependencies. Interdependency exists where one or more software components rely on other components. These components can either be related (part of the same application) or completely unrelated (from a different application). Related or not, there is nevertheless a relationship between the various components. In this relationship, one component requires something from the other. It could also be that both components are co-dependent (dependent on each other) or the dependency might only be one way. Most often, it is functions and API's that are required by one component from the other. However, depending on how the software was written (including the overall design), co-dependent components can and do exist.

To better understand interdependency, it helps to use a simple example. For instance, consider an application and its required software library(ies). In order for the application to work correctly, it requires access to certain functions and API's that can only be satisfied by that specific library. Otherwise, the programmer will have to recreate those functions and API's. This is a very complex and time-consuming process. However, the software library is need not actually be part of the application. Instead, it is from a different application altogether (this is particularly common in UNIX). However, because the programmer wishes to reduce the complexity of his work, using functions and API's developed by others reduces his time, work and expended effort. Programmers tend to prefer building on pre-existing work (e.g. functions, API's, system calls) since doing otherwise can lead to increased programming and operating system complexity.

In large Linux-based operating systems, it is common to have many thousands of such interdependencies. Such a system could have thousands of applications installed and most of them will require functions and API's from any one of the various software libraries found across the operating system. Although it is possible to develop entirely self-contained applications, this is normally not done, as rewriting existing library functionality already available is unproductive.

2.3.3.2 System calls

Another type of interdependency is the system call. A system call, according to [4], is an operating system subroutine that provides a uniform manner for performing operating system-level tasks such as deleting files, managing directories, opening hardware for I/O, communicating with the network, etc. A system call relies on pre-existing functions and API's provided by the kernel. Using system calls, a programmer does not have to recreate a function or API already provided by the system. Furthermore, in all likelihood, the new function would not be as efficiently written. Because the kernel controls the system's hardware, it makes sense to place low-level functionality into the kernel that can be used by developers and programmers alike. When a low-level task (e.g. routine) is called by an application, it places the system call and passes control to the kernel which then runs the appropriate subroutine (e.g. system call) and once completed, returns control back to the requesting application.

2.4 Compatibility issues

2.4.1 Background

Due to its modular nature, the Linux operating system can be replaced in whole or in part. Of course, the issue of software and library interdependencies may need to be resolved when updating or upgrading both software and the operating system.

The openness of Linux makes it possible to see the interdependent nature of applications that rely on various software libraries and the kernel. There are tools that come bundled with most Linux systems that enable the system administrator to examine the various functions and system calls made by requesting applications for library and/or kernel functionality. These tools can also be used to examine library-based issues (e.g. inconsistencies and incompatibilities) that can arise when they are replaced by newer versions. Sometimes, the problem with newer libraries is that older functionality is removed thus causing issues with older dependent applications and libraries.

Nevertheless, this functionality may still be required by certain applications. Therefore, these various troubleshooting tools can also be used to track down specifics that may be indicative of a potential failure when upgrading libraries and or the kernel.

When upgrading a kernel or software library, it is not always possible to determine the impact that changes to these libraries or kernel may have on the system. This is why it is important to test applications after upgrades.

2.4.2 Upgrades

Upgrades impose newer versions of software on the system. Due to the modular nature of Linux, this is generally not an issue. However, existing software that is overwritten by newer software can have far-reaching consequences. These changes could affect a variety of applications and libraries alike. This can introduce interoperability-based issues that are not easily diagnosed or fixed. Furthermore, due to changes that may occur to user preferences and system-wide configurations, it would be prudent to re-examine these after an upgrade. Another issue of concern with upgrades is that certain applications may no longer be supported by either the distribution's maintainer or even by the open source community (in which case support and upgrades are no longer available) or both. In either case, when support for an application is dropped, extra work will be required on the part of the system administrator. If both forms of support are gone, then the application will have to be maintained by the organization itself through various support strategies [2].

Upgrades are generally far more wide reaching in terms of potential impact and consequences than updates. Upgrades generally infer that the system's applications, libraries, kernel and other important subsystems are to be replaced by newer equivalents that may or may not be radically different from their predecessors. Updates, on the other hand, tend to cause minimal system changes that often represent little to no discernible impact for both the users (and their applications) and system administrator.

Upgrades, in the sense used throughout this text, are indicative of an installation program developed by the distribution's maintainer that are intentionally designed and used for upgrading the operating system. Furthermore, upgrades can require a great deal of time to reaccredit and recertify, since upgrades generally impact the entire operating system and not just portions of it, as is the case for most updates. Finally, upgrades are performed for a variety of reasons: sometimes to achieve the next level of functionality or sometimes simply because updates are no longer available for a specific version of an operating system and an upgrade is obligatory in order to continue receiving maintainer-based support as well as provide bugs fixes and patch up vulnerabilities.

2.4.3 Updates

Updates are generally not as far reaching as an upgrade, as the intention of an update is to provide bug fixes, enhanced features, security and functionality. Sometimes updates may also provide enhanced performance and system stability. Nevertheless, the changes imparted onto the system

should not generally affect other programs or libraries[4]. This is because updated libraries are not generally different enough from their predecessors to cause serious compatibility-based issues.

Most if not all updates are transparent to both the user and system administrator in terms of overall system impact. It is rare for system wide and user-based preferences to be changed. Furthermore, unlike upgrades, updates do not generally increase functionality. However, the possibility is always there.

Interestingly, updates can perceptually require as much time as an upgrade to rollout in order for the operating system to be reaccredited and recertified. This is particularly true when many system components are updated on the system. However, with updates, as in contrast to upgrades, specific portions of the system can be updated at any time such as the kernel and its related subcomponents or an application suite.

The key difference between updates and upgrades is that updates generally cease to be available after a certain period. It is common for updates to no be longer available after several years, as the distribution's maintainer will inform customers that support is no longer offered and persuade them to upgrade. Updates, as examined in this text, relate to packages provided by the maintainer for use with an application specifically for updating. Although manual updating is possible, this is generally considered as manual maintenance. Moreover, it is very time consuming and not suggested except as required by circumstances.

Updates are generally used to keep an operating system up to date on an intermittent basis. Sometimes, it may not be desirable to proceed with upgrade just yet. In this event, some updates can be used to keep the system up to date and to benefit from potential feature enhancements in some of the updates. However, this will vary considerably as not all maintainers will provide new functionality in their updates.

2.4.4 Manual maintenance

Performing manual maintenance (via manual patching, recompilation, reinstallation, etc.) is far from an impossible task. It is, however, time-consuming and sometimes daunting, particularly if the system administrator does not have enough experience or if additional resources are needed. Manual system maintenance will require finding recent versions of an application, tool, utility or kernel; and through access to its source code, it is recompiled and reinstalled.

The fact that the Linux operating system is modular makes it possible to install newer applications and libraries without generally compromising system integrity and stability. This in turn enables a more rapid reaccreditation and recertification so that the operating system can be quickly redeployed. In addition, by maintaining vigorous system version control and actively documenting all changes made to the operating system through manual maintenance, it will be easier to successfully upgrade various operating system components, without affecting overall stability. This will further help to increase the efficiency of this process.

[4] However, this will vary from case to case. This is only a general rule of thumb.

2.4.5 Issues

2.4.5.1 Potential problems

Regardless of whether an update, upgrade or manual maintenance is performed, there is always the risk of incompatibilities and/or inconsistencies in the system due to abrupt changes in system libraries. For example, older applications may rely on a set of functions and API's from a given library. However, if that application is no longer supported through a given update or upgrade and the new library is in some way substantially different from its predecessor, the application will very likely become inoperable or too unstable to use.

Of course, there are ways around this, but this can only be determined by attempting the update or upgrade and testing the various applications. If it turns out that the newer library does not provide the required functionality or overly affects system stability, then the previous version can be reloaded from backup media, or using the system's package manager remove the new library and replace it with the previous version. If the new version is nevertheless required, then the system administrator can create a specialized environment where the affected application(s) knows where to correctly locate its required library(ies).

Although this is only an example, it represents an important problem encountered when upgrading or updating an operating system. Thus, the only way to be sure of system integrity after an upgrade or update is to test required and other important system applications.

2.4.5.2 Solutions

It is very likely that the issue of system library inconsistencies can best be managed over a very long-term period, similar to the Royal Navy's requirements of 15 to 25 years, by performing manual system maintenance. Doing this for an operating system is certainly far more complex and time consuming than for updates and upgrades.

The main problem with upgrades, as already stated, is that there will come a point in which older programs will no longer work because (1) they are no longer supported, (2) libraries that they depend on will no longer be supported or (3) those libraries will have been significantly changed will no longer support the older API's and functions. This is never a certainty but it is very likely and the larger a distribution becomes. This is because larger distributions inherently have more interdependencies due to the increased number of programs, tools, utilities and applications.

The main issue with updates is that there comes a point in which the maintainer of the distribution will no longer provide updates for a version of an operating system. As operating systems age and are replaced by newer versions, it is likely that at a certain point, looking to the future, support, updates, upgrades and patches will no longer be available for that older operating system. At this time, it is highly likely that the maintainer will attempt to convince the customer to migrate to a newer and supported version of the operating system, perhaps by offering financial incentives. However, in upgrading the system, a different set of problems will occur over time, as previously explained.

The problem in maintaining the same operating system for such a lengthy period is that there will inevitably come a time when software must move forward to keep up with the times. In most cases, the majority of C2 applications are likely to be custom-built. Therefore, not only the kernel and operating system have to be maintained over the long-term but these C2 applications will also have to be sustained.

Thus, in order to maintain vigorous control over what is to be changed, replaced and how and where newer binaries [5] are to be installed, can only be accomplished through manual system maintenance. While much of the work in manual system maintenance will require the manual reconfiguration, recompilation and reinstallation of programs and libraries, the system administrator has full control over how the reinstallation is to occur and thus exert larger control over how the overall system will be affected. Unfortunately, this requires a system administrator with both ample system administration and programming experience in order to adapt to the various tasks he is likely to encounter. Of course, the system administrator can continue to take full advantage of updates and upgrades in so long as they do not compromise system integrity. *Therefore, it is highly suggested that the system administrator take full advantage of updates and upgrades as long as they are available and do not compromise overall system stability and integrity. Manual maintenance should be left as a final option for maintaining long-term operating system employability*.

The ultimate benefit of performing manual system maintenance is manifest. Through vigorous control of the system and applications, it is possible to minimize or disregard any potential changes to the system. By documenting all system actions and changes, it is possible to realistically trace and narrow down the root cause of problems to previously taken actions or changes. Because of this, it is reasonably expected that reaccreditation and recertification of the operating system will be far less time consuming than with updates and upgrades such that the C2 operating system will be ready for a more rapid redeployment.

However, the need for a trial laboratory in which tests and exercises can be carried out prior to final deployment aboard the frigates cannot be overemphasized. In this lab setting, it becomes possible to determine most potential issues of contention before they become apparent in a theatre of operation.

2.4.5.3 Tools

Most Linux operating systems come with several pre-bundled tools that allow system administrators and programmers to more closely examine the system calls and interdependencies used and required by various applications. The main tools are *Ltrace* and *Strace*. Each of these tools performs a different task, but when combined together, they provide a great deal of information about an application and its dependencies.

Strace can be used to display all kernel-related system calls an application requests. The system's current system calls can be found in various files across the system. Under Fedora Core 6 [6] and

[5] Binaries include runnable executables and compiled libraries.
[6] At the time of this writing in January 2008, Fedora Core 6 was the common Fedora distribution.

17[7], the kernel's available system calls are found in */usr/include/bits/syscall.h*. The *Strace* command can be used to debug many issues that can result from changing a kernel [10, 11]. This command also makes it possible to determine various race conditions that may result from changed system calls as well as the various libraries that are accessed. However, *Strace* cannot be used against actual library files.

Ltrace is another type of program [12]. It is similar to *Strace* but it differs in that it is designed to list all the library functions used by an application. It can also be used to list all of an application's system calls. However, this option is best left to *Strace*, as *Ltrace* does not actually provide the name of the library the function requires nor can it be used against library files.

All running or paused programs, tools, utilities and applications can be queried via */proc/?/maps* and */proc/?/smaps*, where "?" denotes the Process ID (PID) of the application in question [13, 14]. The */proc*-based *map* files can be used to learn information about the various resources a process requires including the memory used by said resources and their permissions. It also lists all libraries in use by said process. The */proc*-based *smaps* files provide highly detailed information concerning the memory resources consumed by a given process. This information can be very helpful in tracking down missing or changed libraries.

Ldd [15, 16] is another tool that can be used to list all the libraries a specific program uses. It is very simple to use and does not require parsing through many pages of information as is the case with *Ltrace* [12]. This program is also able to determine library file dependencies making it very useful for ascertaining other possible library-based contentions. The program can be run against both executables and libraries.

Objdump, another very useful tool [17], is again used to determine which library functions a given program uses. However, the library name is not always stated as the actual file name of the library. Often times, libraries are given proper names and this is recognized and used by *Objdump*. Therefore, combining this tool with *Ldd* [15, 16] will help to sort out any potential confusion. The program can be run against both executables and libraries.

Lesser-used programs that can be of additional assistance include *Readelf* [17] and *Lsof* [18]. It is certain that the abovementioned tools may not be suitable under all circumstances. However, without any debugging information attached to a library or executable, they are an excellent substitute. Furthermore, they are easy to use, even for the novice. More advanced tools are available commercially and are sometimes bundled with specialty development environments for Linux. By combining the information from these various tools, it becomes possible to narrow down many potential problems and once known, fixing them becomes a more straightforward task.

2.5 Summary

Understanding how modern and complex operating systems such as Linux function can be a cumbersome task. Fortunately, it is not necessary to understand all the finer details. However, a basic understanding of dependencies and their effects on the system's overall stability and

[7] At the time this document was upgraded from a technical note to a technical, memorandum in March 2013, Fedora Core 17 was the common Fedora distribution.

integrity is important. Many tools can be used alone or in combination to provide a clearer of view of dependencies and their interaction with one another. The type of system maintenance that will be chosen to keep a system's kernel up to date will depend in part on the expertise of the system administrator as well as the availability of additional resources, not to mention the continued availability of updates and upgrades. There is no clear-cut solution. In certain circumstances, it is more appropriate to apply an update than an upgrade. In other cases, these should be altogether overlooked when manual system maintenance is the simplest and clearest solution for a specific system modification.

Although the Royal Navy has expressed its desire to resist performing operating system updates and/or upgrades, it is very likely that at one point, they will no longer have the choice. It is unrealistic for the Royal Navy to expect that manual system maintenance is the solution for all their maintenance needs. The cost in resources and capital will be too high and complex, especially at the beginning of a deployment. However, the Royal Navy has many options available to it for providing various forms of maintenance [2]. In the end, however, updates tend to be less customizable than upgrades. It is likely that for the first 10 to 15 years, the Royal Navy will not have to worry about going the route of manual system maintenance. Unfortunately, should it occur that a support vendor no longer exists, no longer provides support, has changed its business line or the Royal Navy can no longer approve newer operating system versions due to excessive compatibility-based issues (or other issues), then it will have little choice but to pursue manual system maintenance. Until that time, however, it is not generally suggested.

In conclusion, the Royal Navy will very likely have no choice about performing updates and upgrades if they plan to periodically change hardware. Even a system's hardware that is never changed (which over a 15 to 25 lifespan is very unrealistic) must still be periodically maintained for a variety of reasons, as already put forward earlier in the section. A non-maintained system will eventually suffer from software degradation, a condition that all software systems experience through long-term use and often the only remedy is software maintenance.

3 Methodology I - backup and restoration

3.1 Objective

The objective of this section is to examine in detail the various issues that must be understood before attempting to undertake the task of preparing and creating full system backups. Backups are vital to restoring a system to an operational state should any system changes fail. Failure can occur during laboratory tests, migrations, reconfiguration, updates, upgrades or manual system maintenance. The importance of having a laboratory for testing purposes cannot be overemphasized. In this test environment, various backup and restoration procedures can be examined to ensure the applicability and usefulness of a given backup-restoration scheme, as well as ensure data integrity before proceeding with Part II.

3.2 Backup considerations

3.2.1 Plan development

A plan should be developed prior to performing any backup of the computer systems and their filesystems. This plan should take into account the various facets necessary for a successful backup as examined in the following sections. A backup plan will be the result of a methodological analysis that is a synthesis of the organization's operational policy and requirements that also includes the computers' operating system capabilities and requirements. In addition, a backup plan must be both flexible and adaptable, as contextual changes are likely to occur and must accommodate for unforeseen errors and mistakes. These contentious problems can be introduced by hardware, software and human error. A robust plan will provide for alternate ways around them in order to successfully perform a backup. The backup plan will have a direct impact on the restoration plan (Section 3.4.1) that will itself have to overcome many of the same challenges, thus ensuring the continual availability of data. Finally, as with any plan, it should not be put into action until it has been documented, analysed, reviewed and tested.

3.2.2 Tools

The tools *Dump* and *DD* are well suited for performing Linux and UNIX operating system-based backups. However, other tools such as *Tar*[8] and *Cpio* are more suitable for user and application-based backups. However, different contexts will require different tools for the specific task. The capabilities and ineffectualness of the various tools are examined in the ensuing sections. Although *Dump* and *DD* are preferable for operating system-based backups, *Tar* and *Cpio* have their own specific uses and capabilities. It is assumed that the reader is familiar with the aforementioned tools and understands how to implement them. Finally, no commercial backup tool has been examined in this technical memorandum in order to maintain tool uniformity with respect to the various Linux distributions currently available on the market.

[8] Many Tar compatible archivers have sprung up in recent years.

3.2.3 Use of open source software

Many different software backup tools exist today for UNIX and Linux clients and servers alike. However, the vast majority of these tools are proprietary. The one lesson that must be learned from vendor lock-in is that if a vendor stops supporting a specific product, closes up their business, etc., then the customer or client is stuck with a product that has become obsolete.

In contrast, by using well-known UNIX and Linux backup software tools such as *Tar*, *Cpio*, *Dump* and *Restore*, *DD*, and other lesser used tools, the customer is ensuring that his backups will be useable by not only his site or installation but by any one else using a UNIX or Linux-based system. These aforementioned software backup tools are ubiquitous and are found on all UNIX, Linux and BSD systems. The use of these tools guarantees that backups performed today will be useable 20+ years down the road.

The Royal Navy has many open source backup tools available at its disposition. In addition, various open source frameworks exist that readily make use of the aforementioned backup software tools. The most popular of these frameworks enterprise-ready is *Bacula* that is enterprise ready and supports a wide variety of tape libraries. Of course, the royal Navy is free to choose which backup tool or framework it sees fit. However, the work described in this text assumes that the one or more open source software backup tools will be used.

3.2.4 Use of tape over disk

Today, disk-based backup systems have made large inroads against tape backup systems. However, this does not negate the use or effectiveness of tape-based systems. Modern tape library systems support very large data capacities, with some scaling into petabyte storage. Moreover, tape systems have changed greatly from the time of DDS and DLT based systems. The largest storage standard today is LTO that scale to 6 TB per tape (with 2:1 compression) and with transfer rates comparable to the that of disk. Recent tape libraries not only support LTO technology but also support the latest fibre channel, SCSI, and SAS interfaces for enterprise-based storage scaling. These systems can be taken advantage of using open source backup frameworks such as *Bacula*, thereby ensuring that backups made today will remain compatible 20+ years into the future.

The use of backup technology ultimately resides with the Royal Navy. However, disk-based storage is not particularly well suited to long-term storage, especially when disks are taken offline and placed in storage for the long term [20, 21 and 22]. However, the work described in this text assumes that the a tape-based technology tools will be used.

3.2.5 Data-related factors

An important factor for determining the appropriate backup scheme is the type of data to be backed up. The data type, its location, availability and correlation to the operating system will largely determine which tool or tools can be used. The following subsections are an in-depth examination of these various factors.

3.2.5.1 Data type

A determining factor for selecting the appropriate type of backup scheme will depend on whether the data is operating system, user or application-based. Furthermore, the data type will influence the selection of the backup tool. Operating system data should only be backed up using *Dump* or *DD* (for reasons to be examined further on) while user and application-based data can generally be backed up using any of the four aforementioned tools. These issues are examined in the different subsections below.

3.2.5.2 Special attributes

A determining factor for selecting the appropriate type of backup scheme will depend on whether there are special filesystem attributes that need to be backed up. Backup tools *Tar* and *Cpio* cannot capture special file and filesystem attributes. These attributes can, however, be backed up using *DD* and *Dump*.

3.2.5.3 Devices

A determining factor for selecting the appropriate type of backup scheme will depend on whether there are device files to be backed up. Generally, devices should not be backed up as these are created by the kernel at boot time and are never fixed. Moreover, device nomenclature has changed over the years and will likely continue to evolve. Using tools such as *Tar* and *Cpio* to backup devices can cause these tools to actually attempt to read from these devices, thereby backing up all disks and other associated devices. If devices must be backed up, use *Dump* or *DD*.

3.2.5.4 Raw data

A determining factor for selecting the appropriate type of backup scheme will depend on whether any raw partitions or disks need to be backed up (e.g. raw database partitions). *DD* is uniquely suited to this task. *Dump*, *Tar* and *Cpio* cannot be used because there is no recognizable filesystem for data acquisition.

3.2.5.5 Locked files

A determining factor for selecting the appropriate type of backup scheme will depend on whether locked files need to be backed up. Certain key operating system files from */dev* and */proc* are automatically locked by the kernel and are normally not accessible by most of the backup tools. These files are also generally off limits to users and applications (and sometimes the root user). Consider the following:

a. Are there locked operating system files? *DD* is generally the only way to acquire them if the *Dump* program cannot otherwise acquire them. *DD* should only be used against a filesystem when the partition or disk is offline or mounted read-only.

b. Locked pseudo-data files found within */dev* and */proc* can generally be acquired using *Tar* and *Cpio*. Know exactly which files and pseudo-files to acquire beforehand.

c. Consider what files and pseudo-files are required and for what reason. Consider that reading from special files such as */dev/sd??*, */dev/mem*, */dev/kmem*, */proc/kcore*, and many others will cause the backup program to read from system reserved devices that will:

   i. System instability;

   ii. Attempt to acquire all device-associated data therein. For example, reading from */dev/sda* will acquire all disk data associated with that device while reading from */proc/kcore* will effectively dump the system's memory.

   iii. *Dump* generally does not copy the data therein as these directories, data files and pseudo-files are based on the currently running kernel.

   iv. */proc* and */dev* are best left alone except in the case of a full filesystem dump where *Dump* handles them without issue and *DD* which acquires the underlying partition or disk.

d. Are there user or application-based locked files? Can those files be backed up without resulting in file or data corruption? Generally, other than for database files, locked application and user files can be backed up using *Tar*, *Cpio* and *Dump*. However, locked operating system files cannot.

e. If some files must be backed up and are locked and inaccessible no matter what is attempted, then booting from a rescue or Live CD or other operating system disk can be used to back up these files.

3.2.5.6 Data file volatility

A determining factor for selecting the appropriate type of backup scheme will depend on whether the files to be backed up contain volatile state-based information about the operating system or some arbitrary application or service. Consider the following:

a. Operating system files that contain volatile data from directories such as */dev* and */proc* should **not** backed up. Many of these files, especially those from */proc*, contain highly volatile system and application data (e.g. RAM, system and application states, etc.) and must be treated carefully.

b. Many application and service-based and can be safely backed up, even when the files are in use. *Dump*, *Tar* and *Cpio* can normally handle this without issue.

c. Backing up device files (found under */dev*) can be problematic. Usually, device files do not need to be backed up. However, a full filesystem dump of the operating system will nevertheless backup these files. *Dump* can normally backup these device-based files even if they are in use.

d. If for some reason certain files that must be backed up cannot be accessed regardless of various attempts to back it up, then the system must be placed offline and backed up using a rescue or Live CD or booting from other operating system disk.

3.2.5.7 Running applications and services

A determining factor for selecting the appropriate type of backup scheme will depend on whether application-based files can be backed up while the application is left running. For example, it is generally not advisable to back up a database while the database service is running. Buffered data may not have yet been written out to the database, potentially resulting in data loss upon restoration. Consider the following:

a. An application or service left running can normally be backed up as is. *Tar* and *Cpio* are well suited for application-based backups. *Dump* can also be used.

b. Where database applications are concerned, the application or service and its associated files (configuration files, binary executables, libraries, etc.) can normally be backed up. However, databases, in of themselves, should never be backed up while in use. The database application or service should first be stopped.

c. Many small database files can be backed up while their corresponding applications are running. This however, will depend on the type and size of the database. Larger databases and should always be shut down prior to backing up. This is due unwritten buffered data in memory. For smaller databases, unwritten buffered data may result in nothing more than several minutes' worth of data loss. Larger databases, however, may be inconsistent and be rendered unusable if buffered data is left unsynchronized.

d. Is the application data stored on formatted or raw devices? *Tar*, *Cpio* and *Dump* will function against partitioned devices where valid filesystem structures exist. *DD* is best suited to raw devices.

e. Are the applications services? They too can normally be backed up. However, services often write data out to log files and these log files may not have been recently synchronized such that when they are backed up, they might not contain memory-resident buffers, possibly resulting in partial data loss. If a service's log file(s) is important, then the corresponding service should be shutdown prior to backing it up. Normally, backing up services and their various files is not an issue.

f. Are user data files currently in use by applications? Normally, user data files (e.g. word processing, spreadsheets, etc.) that are in use can be backed up. However, much will depend on file sizes and how often buffered data is synchronized. Normally, this will not result in corrupted files but may result in partial data loss. Therefore, user files in use can be backed up using *Tar*, *Cpio* and *Dump*.

g. If certain files are locked and the corresponding application or service cannot be shut down, then certain files may have to be skipped. If the filesystem is active but quiescent, then *DD* can be used to acquire an image or the system can be taken offline and backed up from a rescue or Live CD or other operating system disk.

h. Are files being shared over the network? These files can also be backed up. However, if they are being edited or modified, then what is in memory may not be consistent with what was backed up, unless unwritten buffered data is synchronized out to disk before backing them up.

In summary, it should be considered a good practice to shut down all non-essential user-applications and databases before backing them up. If they must be left running, ensure wherever possible that unwritten buffered writes are synchronized to disk. *Tar*, *Cpio* and *Dump* can be used liberally for these types of backups.

3.2.5.8 Active operating system

A determining factor for selecting the appropriate type of backup scheme will depend on whether the operating system must be left running while backing it up. This will depend on a host of issues such as file types, special file locking mechanisms and the availability of devices and files. Generally, most if not all operating system files can be backed up while the system is running. However, it may be necessary to shut down certain services and user-based applications before proceeding with a backup. Consider the following:

a. What is the availability of the operating system's files? Are they all available? Are some unavailable? Can the backup proceed without those files? Generally, once applications and services have been shut down (only if necessary) the only inaccessible files are volatile operating system and device-locked files that are not generally required for operating system backups. However, if one or more of these files must be backed up, then an offline system backup using *Dump* or *DD* can be carried out using a rescue or Live CD or other operating system disk.

b. If necessary, can the operating system be shut down for a given period to back it up? This may depend on organizational policies and operational requirements. An inactive operating system and its files can always be backed up when it is offline, regardless of file type, data contents, access lists, permissions, etc.

c. Can portions (e.g. partitions) of the operating system be taken offline or made read-only in order to decrease backup time and system downtime.

d. Is it important to back up the more volatile portions of the operating system such as */proc*? Generally, the answer is no as most of its contents are volatile and dynamic and is only required by the running kernel. Only certain static text configurations are worth backing up from here, but these configurations should always be initiated from start-up and initialization files and scripts. However, if necessary the *Dump* command can be used to back up certain contents of */proc* without taking the system offline.

e. If certain files cannot be backed up while the system is online, will a restoration be successful without those files? If not, are there other available sources for those files? If not, then the operating system will have to be taken offline in order to backup those files.

f. Is the operating system spread across multiple disks and/or partitions? If so, then often these operating systems are easier to backup than single-partition systems. There are

many reasons why this type of system is easier to backup. *Dump* can be used for almost all the partitions and where it cannot be used, if a given partition can be brought offline or is inactive, and then *DD* can be used.

g. Many other factors exist that must be contended with that go beyond the scope of this subsection. Whenever system volatility and file accessibility -based issues are present, backup procedures will generally be more cumbersome than they would otherwise have been. Thus, there is unfortunately no ubiquitous solution for operating system backups, only practical solutions.

3.2.5.9 Filesystem availability

A determining factor for selecting the appropriate type of backup scheme will depend on whether filesystem(s) can be placed offline. This may vary according to various system requirements and organizational policy. Consider the following:

a. Can the files be safely backed up if the filesystems are mounted? If so, then *Tar*, *Cpio* and *Dump* can be used. *DD* should never be used to acquire actively mounted partitions unless the filesystem(s) is quiescent.

b. If one or more partitions are placed in read-only mode, can the operating system carry on for the duration of the backup without affecting system stability? If so, then *Tar*, *Cpio*, *Dump* or *DD* can be used.

c. Will placing filesystems in read-only mode affect applications, services and user data files? If so, then either backups should be performed during off-hours or if this is not possible, disable the affected services and applications.

d. If one or more partitions are placed offline, can the operating system carry on for the duration of the backup without affecting system stability? If so, then *Dump*, *Tar* or *Cpio* can be used.

e. Are locked files preventing filesystems from going offline or being made read-only? If so, then by disabling the service or application using the locked file(s) may correct the problem. If it does not, then the system itself may have to be placed offline in order to obtain a backup of the filesystem.

3.2.6 Other backup factors

3.2.6.1 Media

A determining factor for selecting the appropriate type of backup scheme will depend on the type of backup media that will be used, its location, its path and its availability. The backup media can be local or remote to the system to be backed up and it can be a disk or tape. It is important that media access always be available under varying conditions. For example, if data is normally backed up to a local tape drive and it fails, then a remote tape drive can be used but only if the

network path is available and permissions are set for sending and receiving a remote data stream. Consider the following:

a. Is enough media available, especially if tape is used? Is the backup tool multi-volume capable? *DD* is not multi-volume capable, but it can be made so through scripts.

b. Can the media be read and written to by other devices? For example, is the media universal enough that other devices (e.g. tape drive) can read and/or write to/from it?

c. Are spare local and remote backup media and devices available to parallelize the backup if multiple disks and/or partitions can be simultaneously archived? The backup tool does not need to be multi-volume capable in this case, only multi-device capable (achieved through scripts).

d. If the media and/or archival device are remotely located, are the necessary remote system and network permissions in place for data streaming and reception?

e. Are the device, media and network paths available? Are there specific schedules for their usage or are they available 24x7? Can they be rededicated for other tasks?

f. Is the media checked for errors prior to usage? If not, can this be automated using scripts? Does the backup device perform error or CRC-checking?

g. Are different types of backups stored on different media? For example, are backups done at different times placed on the same or different media? If an incremental and differential backup of one or more systems is done, is it also placed on the same media?

h. Are different types of data stored on different media? For example, will operating system backups be placed on the same media as user data backups?

3.2.6.2 Lifespan and storage

It is important to determine how long data should be kept. Depending on the data and its function, it may be necessary to keep it for several weeks, months or longer. For example, legal documents must be kept by law for a prescribed period. Consider the following:

a. Is the data to be kept for a short or long period? Are there any laws that mandate how long certain types of documents must be kept? What is the organizational policy concerning backup and data longevity? Is the policy the same for user data and operating system data?

b. If the data is stored for a long period, there are various storage issues and requirements to examine. For example, will the media be stored in a temperature and humidity controlled environment? Is it secure, safe from fires, floods and tampering? Who has access to it? Will it be stored onsite or offsite?

c. Will data be backed up incrementally or differentially? This may depend on required data longevity and organizational policy. Is there a preference for incremental or

differential backups? All the tools but *DD* can be used for incremental or differential backups. How often is a full system backup (operating system and user data) carried out with respect to incremental and/or differential backups? These answers will help to determine how much media is required.

d. Are the users involved in helping determine the lifespan and usefulness of their data? Is there a policy in place for maintaining user data backups for set periods? For example, should users be consulted on an intermittent basis to determine if their data and storage requirements have changed?

e. What is done with users and their data when they are no longer system users? Are the accounts destroyed and all data saved to CD or DVD or must the accounts be disabled and data stored and backed up with all the rest of the user data?

f. How often are full backups performed? Are they incremental or differential? Once a new full backup is performed, existing incremental or differential backups could be overwritten and reused again for newer incremental or differential backups.

g. Modern tapes can be used for at least several thousand uses. They can be left in a drawer or cabinet for long periods and still be expected to work; assuming proper long-term storage conditions were met. Disks, on the other hand, cannot be left idle and not spinning for very long periods as they tendency to break down increases the long they are left unused. [20, 21 and 22]

3.2.6.3 Data security requirements

Before carrying out any backup, it is important to determine the security requirements necessary for safeguarding data, both during backup and storage. Consider the following:

a. Will data be sent over the network? If so, is it on a trusted network? If not, then perhaps the data should be encrypted. This can be done by piping the backup data stream through *SSH* or *GPG* before it goes out onto the network.

b. Should the data be encrypted, regardless if it is backed up locally or remotely? If so, then the data should be encrypted from the source system using *GPG* or other similar tool.

c. If encryption is used, is there a mechanism or resource to help in the management of encryption keys (e.g. PKI)? Who has access to the keys? Who has the password s necessary for encryption and decryption?

d. What network tool should be used for network backups? *RSH*, *SSH* and *Netcat* can be used. *SSH* should be used for encrypting the network stream between computer systems. *RSH* and *Netcat* can be used on trusted networks or if the data stream is already encrypted via other means (e.g. *GPG*).

e. What are the current security settings for the machines involved in the backup process (local sources and remote systems)? What should the security settings be set to? Can

unauthorized users or systems gain access to either the source and/or remote systems while backups are being performed?

f. Does the backup tool support I/O piping? This may be required, depending on the type of backup to be done, including compression and encryption requirements as well as data backup location (local or remote backup system). *Tar*, *Cpio*, *Dump* and *DD* all support piping. Advanced piping capability can also be achieved through scripts.

g. Does the backup user have the necessary rights and permissions to backup all system and user data? Does that user have permissions to send and receive backup data streams over the network to and from remote systems?

3.2.6.4 Size requirements

It is important to determine the various size requirements of the data before performing any backup. Consider the following:

a. Is there enough backup media available for the backup to complete correctly?

b. Is the backup going to require multiple volumes? If the backup cannot be stored onto a single backup media, will the backup tool require multi-volume capability? *DD* can be made multi-volume capable using scripts.

c. Is backup space a concern? Does the backup-based device support hardware compression? If space is a concern and the backup device does not support hardware compression, then through piping, any of the backup tools can be made to send the backup stream through *Gzip* or *Bzip2* for compression before redirecting the data stream to the backup device.

   i. Modern computer systems equipped with many multi-core CPUs can take advantage of recent developments in parallel compression. Highly capable parallel compression software includes *pbzip2*, *pigz*, *pxz* and many others.

d. Large datasets such as operating systems or databases should be backed up using *Dump* because it is much faster than the other tools.

e. Small datasets and user data can be easily backed up using *Cpio* or *Tar*. However, larger datasets should use *Dump* wherever possible.

f. Modern tape drive and media now scale at up to 6 TB per media unit and tape libraries scale to well over petabyte storage capacities.

3.2.6.5 Data accuracy and relevancy

It is important to determine how accurate the data to be backed up must be prior to performing any backup. Consider the following:

a. Is the exactness of the data important? Must it be exactly as it appeared on disk? If so then *DD* should be used to acquire a bit-copy image.

b. Are raw partitions to be backed up? If so, raw partitions must be bit-copied to ensure accuracy, as there is no recognizable filesystem on them. *DD* should be used for such a situation.

c. Are the filesystems too large to be backed up to one media? If so, then the backup should be done using a multi-volume capable tool or script.

d. Does the backup device support CRC or error checking when reading or writing data? Can the backup tool perform CRC or error checking? Can it deal with errors and if so, how? How much data will be lost if bad blocks are found on the media?

e. Does the backed up data contain the same filesystem attributes as the original data? Does the tool understand how to handle extended filesystem attributes?

f. Can the backup tool accept user-specified input for determining which files to back up?

g. If disk exactness is not necessary, then the tools *Dump*, *Tar* or *Cpio* can be used.

3.2.6.6 Speed and bandwidth

It is important to determine the network speed, the amount of time available and required by the backup tool, including the required network bandwidth. Time constraints will also be a determining factor in deciding which backup utility to use. Consider the following:

a. *Tar* and *Cpio* are very fast when backing up small to medium-sized datasets. *Dump* is very fast for datasets of all sizes. *DD* is not fast because it must copy every bit of data from a partition or raw device.

b. How much network bandwidth is available? Backups should be scheduled when network bandwidth is at highest.

c. If backups must be performed during periods of high network utilisation, then data compression tools can be used to help reduce network bandwidth utilisation. The data stream can be compressed using *SSH*, *Gzip*, *Zip*, *Bzip2*, *pbzip2*, *pigz*, etc., before being sent and decompressed at the remote system.

d. If *DD* must be used, then it will typically require more network transmission time as bit-copies are larger than their filesystem-based counterparts. *DD*-based backups should be compressed prior to network transmission and backed up during periods of low network utilisation.

e. What is the read/write speed of the archival media and/or device? This could be the bottleneck. Spreading out the backup from multiple partitions or filesystems across multiple backup devices (local or remote) could significantly speed up the backup.

However, this is not a useful option for multi-volume backups. Instead, it is better suited for simultaneously backing up multiple independent datasets.

f. Local backup devices are usually faster than remote devices and using multiple backup devices simultaneously for large datasets from different partitions is faster than using a single backup device, whether local or remote. However, the local system must support the required I/O bandwidth.

g. SCSI[9] devices are generally faster than IDE[10] devices and are therefore highly suggested.

h. Can the network scale to the required bandwidth necessary for backups, especially if multiple network backup streams will be in use concurrently?

3.2.6.7 Tool summary

It is important to determine the capabilities of the various tools examined within this section. The following provides a brief summary for these tools:

a. Will data be sent over the network for backing up to a remote system? If so, then *RSH*, *SSH* and *Netcat* can be used for the transmission and receiving of network-based backups and restorations.

b. Is the backup to be multi-volume? If so, then *Tar*, *Cpio* and *Dump* can be used. *DD* can only be made multi-volume capable through scripts.

c. Is data encryption required? If so, then *GPG* or another similar tool can be used.

d. Is network encryption required? If so, then *GPG* or *SSH* or another similar tool can be used.

e. Is data compression required? If so, then *Gzip*, *Zip* or *Bzip2*, *pbzip2*, *pigz*, etc., can be used.

f. Is network compression required? If so, then *SSH*, *Gzip*, *Zip* or *Bzip2*, *pbzip2*, *pigz*, etc., can be used.

g. Is piping required? If so, then *Tar*, *Cpio*, *Dump* and *DD* can be piped into all of these programs: *RSH*, *Zip*, *SSH*, *Netcat*, *GPG*, *Gzip* and *Bzip2*, *pbzip2*, *pigz*, etc.,

h. Is error checking required? If so, then *Tar*, *Cpio* and *Dump* support some form of error checking. *DD* can perform error checking using scripts.

[9] SCSI devices are no longer in high use in 2013. While they still exist, they have been largely superseded by SAS-based technology.
[10] IDE drives have been largely superseded by SATA-based technologies permitting larger data densities and transfer rates.

i. Is the operating system to be backed up? If so, *Dump* or *DD* should be used, depending on the backup requirements of exactness.

j. Are user and application data to be backed up? If so, then *Tar*, *Cpio* or *Dump* can be used.

k. Is the dataset large? Then *Dump* or *DD* should be used, depending on the backup requirements of exactness.

l. Is the dataset small? Then *Tar* or *Cpio* should be used.

m. Are incremental or differential backup required? If so, then *Tar*, *Cpio* or *Dump* can be used.

n. Are there locked files? Are there volatile files? If so, then *Dump* should be able to handle them.

o. Are applications and user data in use? *Tar*, *Cpio* or *Dump* should be able to handle this.

3.2.6.8 Resource availability

It is important to determine the availability of various backup resources prior to performing any backup-related action. This will help to better plan and schedule resources. Consider the following:

a. Are centralized backup servers available? If so, a backup server can be used to centralize and facilitate backups from one or more systems. However, the more systems that stream data to a centralized system at the same time, the more network bandwidth will deteriorate.

b. Does the centralized system have more than one backup device or library? If so, then more than one system can back up its data at the same time. However, performing multiple backups and/or restorations require more system I/O and network bandwidth.

c. Is the backup to occur over the network onto a remote system? If so, can the network support the backup system's required network bandwidth? If not, then multiple network adapters with different network connections will be needed to handle the network load.

d. Are spare backup systems and devices always available? This is important if a backup fails due to hardware. Having additional remote systems (if they are used) and extra backup devices will facilitate and speed up the backup process should a hardware failure occur.

e. Is there enough spare media? Every so often, media is found to be defective, even new media although for tape this is very rare.

f. Is the amount of data to be backed up very large? If so, then it may be more appropriate to backup the system's various filesystems simultaneously onto multiple backup devices (local and/or remote devices) to speed up the process. Consider using a tape library.

g. Are backup operators required or available to intervene if backups should fail or require a specific action to be taken?

h. Are the archival devices capable of handling multi-volume backups? Robotic tape libraries can handle multi-volume backups and can be managed using backup frameworks including *Bacula*. If the device is not a tape library, then a multi-volume backup can be performed using scripts that divide a backup among multiple backup devices.

3.2.6.9 Backup schedules

It is important to determine when backups can be scheduled. Backups should be carried out prior to implementing any noteworthy system change or modification. Consider the following:

a. When are the systems to be backed up? Is it to be done at night, during the day or on weekends? Backups should be done when the necessary resources are available.

b. Are network resources available 24x7 for remote backups? Is there more bandwidth available on weekends and at night? If so, then backups should be scheduled for these times.

c. Are backups planned to coincide with increased network bandwidth availability?

d. Will in-use applications and user files affect backups? If so, then backups of user and applications files should occur when users have logged off. However, if in use files and applications do not affect backups, then they can be scheduled for any time.

e. Do certain operating system services and/or applications interfere with the backup process? If so, then a schedule should be made when these service or applications can be disabled. If they do not interfere, then the backup can proceed at any time.

f. Are certain data files locked or unavailable? If the backup can continue without them, then the backup should be carried out. However, if these files are required for a successful backup and the tools *Dump* or *DD* are not able to acquire them, then a time should be scheduled to place the system offline for backup using a rescue or Live CD or other operating system disk.

g. How will backups be scheduled and executed? If they are using the system scheduler (e.g. *Cron*, *At*, etc.), then backups can be automated. If they are executed from the command line, then a backup operator will be required. In either instance, scripts can be used for automation.

3.3 Filesystem checking for backups and data restoration

3.3.1 A note about disks and RAID arrays

Most UNIX and enterprise Linux systems use non-arrayed disks for the operating system and related data. RAID arrays are typically used to store user data, databases and large datasets. Thus, it is assumed that in this subsection, unless otherwise stated, that when dealing with a disk or its filesystem both individual and arrayed disks are being referenced. Although Linux supports software RAID arrays, they are never suggested in lieu of more robust hardware RAID controllers.

3.3.2 Filesystems

All filesystems, no matter how robust, modern or technologically perceptive, are all at risk for data loss and/or corruption. The most likely time for this to occur is when memory cache buffers have not been synched to the filesystem. Although modern filesystems are able to reduce this danger by incorporating filesystem journaling, the risk is always present. In addition, unrecoverable hardware errors and kernel panics [11] may also be likely culprits. This section therefore examines various filesystem-checking tasks that should be performed prior to undertaking any backup-related action.

3.3.3 Reasons to conduct filesystem checking

Filesystem checking, although very important for backups, are equally important for data restoration. Restoring files from a known "good" backup to a damaged filesystem could result in the overwriting of existing files or even further corrupt the file system. Similarly, a corrupt backup due to filesystem damage will likely result in missing files upon restoration, assuming that the archive is not corrupt because of the damaged filesystem. Some readers may assume that these issues are exaggerated. However, a simple lookup on the web for filesystem errors with respect to backups and restoration will yield sufficient horror stories.

3.3.4 Periodic checks

Filesystems must be checked periodically. Even the simplest filesystem inconsistency can cause data loss or damage. Moreover, severely damaged filesystems may be missing files or may become inaccessible due to damaged permissions and/or ACL's. The only way to verify the status of a filesystem and its files is to perform a filesystem check. Linux provides the necessary tools for all of its supported filesystem formats, the most notable of which is *Fsck*. However, certain filesystems are more sensitive than others are to damage. For example, Ext2 is far more sensitive to power outages and incomplete data writes as compared to its successors, the Ext3 and Ext4 filesystems. Both filesystems employ a journaling mechanism to ensure safe data writes resulting in a more consistent filesystem.

[11] A kernel panic is generally triggered by an irrecoverable hardware or software error.

In addition, certain backup tools are more sensitive to filesystem errors than others are. *Dump* and *Restore* are particularly sensitive to filesystem data structures. These data structures are stored using filesystem inodes and damage to them can result in damaged or inconsistent backups. It is also possible for backups to crash and result in incomplete archives due to filesystem errors. This can occur for all the backup tools examined herein. Nevertheless, all these tools can suffer from filesystem inconsistencies.

3.3.5 Repairs

The Linux operating system supports and provides all the necessary tools for checking and repairing a variety of filesystem formats where all but the most troublesome errors can be fixed. Filesystem checks and repairs are carried out using the *Fsck* utility. Although filesystem checking is important, there are certain dos and don'ts. The most important don't is that a filesystem should never be repaired when it is mounted, even if it is mounted in read-only mode. Even when a read-only filesystem is mounted, buffered changes made to it may result in filesystem inconsistencies. File systems can be checked, **but not repaired**, when they are mounted, although this will vary according to the filesystem in use and may be dangerous. It is always safest to check and repair a filesystem when it is offline. *DD* disk image files can also be checked and repaired just as with any other valid partition or filesystem and the same basic rules about checking and repairing apply. Severely damaged filesystems (if a backup is not available) may require forensic recovery to reconstruct damaged data and/or filesystem structures. Finally, the root filesystem should always be checked and repaired in single-user mode (during system start-up) or from a rescue or Live CD or other operating system disk. Under no circumstances should the root filesystem be repaired while the operating system is fully functional and operational (e.g. Linux runlevel 3 or 5).

3.3.6 Scheduling

It is important that time be made periodically for checking a system's filesystem(s). Often, the best time to do this is at system start-up. However, some systems, particularly mission critical C2 systems, may have operational requirements that make taking them offline difficult. Nevertheless, periodic maintenance is required even on these systems. If time cannot be periodically scheduled for certain systems, then filesystem checks and repairs should be made before performing other system maintenance tasks such as updating, upgrading, and/or replacing/adding hardware. Although certain filesystem formats are almost "impossible" to corrupt (e.g. XFS) due to the many inherent safeguards they incorporate, they should still be checked. There is no such thing as a perfectly stable filesystem, since they are always changing and in a constant state of flux. In addition, if it has not already been done, incorporate filesystem checking into the organizational policy regarding system maintenance. Finally, it is better to have certain systems or services unavailable for brief periods, rather than have an inconsistent or corrupt backup or restoration when it really matters.

3.3.7 Filesystem formats

Linux supports multiple filesystems formats and provides the necessary tools for fixing most of them. However, how well a given tool can support, check and fix a filesystem will largely

depend on the maturity of the filesystem and the support afforded to it by Linux. For example, the UNIX UFS/FFS filesystems are readily used by Solaris and BSD; while they are supported by Linux, there are no filesystem checking tools available for them. NTFS filesystems are also very popular for Linux systems but again no filesystem checking tools are available for it under Linux.

All of the major filesystems commonly used by Linux are fully supported. In addition, Linux supports the mounting of loopback virtual filesystems (e.g. *DD* disk image files, ISO files, etc.).

3.3.8 Bad blocks

Although much attention has been paid to the file system, it is also important to check a disk's surface for physical errors. Disks can suffer from excessive wear and tear due to the movement of disk read/write heads. This wear and tear can result in the formation of bad blocks. These bad blocks are damaged physical disk sectors that contain actual information. However, that information may or may not be useful just as it may or may not reside in unused disk space. What information is contained within a disk block will depend on where it resides with respect to the underlying filesystem. Although a few bad blocks by themselves will not result in an inconsistent filesystem, they can cause files to become damaged that may lead to data loss or corruption. The Linux operating system provides two tools to check for bad blocks and move damaged filesystem blocks to "good" spare blocks. These are the *Badblocks* and *Smartmontools* disk-checking programs.

It is important to perform bad block checking before any errors become disruptive to the operating system or users. Bad block checking as per the *Badblocks* program can be carried out against a mounted filesystem so long as a non-destructive read-only check is used. Destructive or read/write checks must be conducted against unmounted filesystems. However, for C2 systems, taking filesystems offline is problematic and may require appropriate resource scheduling. The *Badblocks* program is useful for two purposes. The first is as a general disk-checking tool and the second is to pass a list of detected bad blocks to the filesystem formatting command (*mkfs*). *Badblocks* does not recover data from detected bad blocks.

The *Smartmontools* can be used to check disks and to move data from detected bad blocks to spare disk block. Although *Smartmontools* has some data recovery capabilities, they are limited. If a disk has been detected as having an abnormal number of bad blocks that disk and all the filesystems contained therein should be placed in read-only.

Thus, if no recent backups have been carried out against these filesystems, perform these backup as soon as possible and then take the disk offline and replace it. If the disk is contained within a RAID array, then the RAID hardware controller can generally rebuild and reintegrate the disk into the array.

3.4 Restoration considerations

3.4.1 Plan development

A successful restoration plan should be based largely on the backup plan (developed in Section 3.2.1) that should already have examined the different facets involved in backing up Linux-based computer systems. These facets include organizational and operational policy and requirements, as well as the operating system capabilities and requirements. How a restoration unfolds is largely dependent on how the backup itself is done. This can include backup devices, the use of centralized backup servers, the type of media and filesystems, skipped files, network paths, etc. Restorations are important to recover from data loss, corruption or operating system error. However, an improperly backed up system may cause restoration failure and the files may be permanently lost. Consider that if a particular backup is difficult to perform due to various factors, then it is more likely that the restoration will also be difficult to perform. Conversely, a backup done with ease should be also easily restored. A restoration, as with any backup, can go awry for many reasons, the majority of which can be mitigated or altogether removed through adequately planning. This subsection will help to plan for a successful restoration.

3.4.2 Tools

The tools used for data restoration are the same ones used for performing data backups. *Tar* is fully capable of creating and archiving data into a *Tar*-based archive. It can also extract data from the archive. It is the same for *Cpio*-based archives. *Dump*-based archives are restored using the *Restore* tool. It is a command-line and interactive tool used for extracting files from *Dump*-based archives. Finally, *DD*-based can be restored in one of two ways. The first is to restore the bit-copy disk image to the same disk from whence it came from or to another disk of compatible disk geometry. The other option is to mount the disk image using the Linux system's virtual loopback feature and restore the filesystem data (or parts therein) to the target disk where the restoration is destined for.

3.4.3 Various restoration factors

3.4.3.1 Resource allocation and assurance

It is important prior to any restoration that all required resources be made available and allocated. It is important that resources already be available before the restoration is to commence, as unforeseen delays will inevitably lead to increased down time for systems, and mission critical systems can only be taken offline for short periods.

Network paths, centralized servers, remote backup devices, etc., should also be available when required and tested beforehand to determine their functional status. In addition, spare hardware should always be on hand in order to mitigate and resolve hardware-based problems as they arise.

Finally, it is important to have a rescue or Live CD or other operating system disk to boot from in the event that the restoration of the operating system fails so that additional restorative attempts can be made.

3.4.3.2 Scripts

The restoration scripts, if any, should perform the inverse function of the backup script(s).

3.4.3.3 Data safeguarding

All backup media should be safeguarded and stored in cool environments as stated by the manufacturer's long-term storage recommendations. Generally, this will require controlled environments that have the appropriate mixture of temperature, humidity and stored away from strong magnetic sources. In addition, the media should be stored according to its value. High-value data should be stored nearby but off-site, possibly in fireproof or fire-resistant storage containers, lockers or safes. Furthermore, the media should always be secured from theft and tampering. It is also important to consider placing spare equipment such as backup devices and disks in lockup, so that when and if they are required, it is certain that they will be there.

Safeguarding offsite is not possible when backups are made aboard a frigate. In such cases, it best that the media reside elsewhere aboard in a safe and temperature-controlled environment.

3.4.3.4 Media testing and device diagnostics

All backup media should be tested periodically to ensure an error free state. In so doing, this will enable the media to be used at any time without worrying about added complications during data restoration. However, exactly how the media should be tested will depend on a variety of factors. For example, the type of media used and its archive format will be important determining factors. In addition, various tools can be used. *DD* can always be used on tapes to perform tape dumps, regardless of the underlying archive format. The archiving tool, of course, can always be used on its respective archive. If the archives are stored on disk, then the same process of testing the archives continue to apply.

Obviously, the more media there is to test, the longer testing will take. Testing time, however, can be reduced by using scripts and this can be amplified by using automated tape libraries that can be controlled by a script program or through a more advanced framework (e.g. *Bacula*). In addition, disk checking and recovery can be automated using scripts. However, an important question to answer is how often should the media be tested? Another is should the entire media be tested or individual archives? Often, testing once per year will suffice. Testing can often be done by dumping a table of contents of the archives. This is the default for most tools, although full data extraction is always a possibility. If tapes are used as the main form of media, then they should be re-tensioned periodically to ensure that the tapes do not stick.

It is also important to test the tape drives, disk controllers and disks (if media is on magnetic disk) using standardized procedures and best practices. Often, the devices' manufacturers may provide useful insight and tips. It is important to implement an organizational policy if one is not already in place.

3.4.3.5 Alternate methods of restoration

It is important to ensure, as with backups, that the restoration plan examines and outlines various alternative methods for restoring data back to the required system(s). As with backups, it can happen that backup devices or network paths become unavailable and another method for restoring the system(s) must be found. Unforeseen and uncontrollable circumstances can cause data restorations to become very complex and cumbersome when being performed against mission critical systems. For example, if certain disks or filesystems cannot be brought offline on the system marked for restoration, then data can be restored to an alternative directory (on the same system or on the network somewhere else and then shared) and then copied back over at another time. Alternatively, if network restorations are performed but the network becomes inoperable, then it is important that the backup device and media can be easily connected and reconfigured for the system that is to have its data restored.

3.4.3.6 Procedure testing

It is very important to test all restoration procedures before ever requiring them. Testing should be done in a laboratory or other controlled environment where different tests and scenarios can be experimented on without affecting the operational network. Procedures should be tested to determine their suitability and adaptability to changing contextual situations in order to understand how to improve upon them, when and where necessary, so that they can be made to accommodate additional uses. In addition, by testing them, it becomes possible to understand how they will behave when faced with to varying circumstances and other mitigating factors.

However, not all procedures are equally useful nor are they equally applicable in all circumstances. For instance, backed up system data that is stored on tape that is no longer accessible due to a failed local SCSI controller will have to be restored over the network. Another example would be a system's main operating system and boot disk that has failed. It will have to be replaced with a spare disk and its data reloaded from backup. This will require booting from a rescue or Live CD or other operating system disk. In the latter example, procedures should already be in place for how to use alternate boot devices and how to restore data using these devices. Regardless of the circumstances, it is important that the backup media used can be recovered using the procedures in place which should be found in the restoration plan. The most difficult issue concerning a restoration is adapting it to varying circumstances when time is running out, as for mission critical systems. Thus, a set of procedures will help to better adjust to such circumstances.

3.4.3.7 System and filesystem availability

Depending on the type of data to be restored, it may be possible to restore the data directly back to the affected filesystem(s) without interfering with the underlying system's ongoing operations. However, this will depend greatly on the type of data to be restored. Generally, applications, services and user data can be restored without ever having to take the system offline. Certain applications or services may require being taken offline temporarily while their data are restored. In addition, depending on the type and amount of user data to be restored, it may be necessary for users to logoff while their data is restored. It is also possible that all the users may have to logoff, depending on how much user data is to be restored and whether it will affect some, most or all of

them. If the backup contains operating system data however, then it is possible that the system will have to be brought offline in order to conduct such a restoration. This too will depend on whether the backup must overwrite system-locked files or other critical operating system files that must not be changed while the system is operational. Thus, a determining factor for taking a system offline will depend on the type of data to be restored.

If a backup was conducted using *DD* then the disk image can be restored directly to the affected disk, but only if the target disk's filesystem(s) is not mounted. Otherwise, the disk image can be logically mounted using the Linux system's virtual loopback device and the filesystems therein restored using a variety of tools and techniques. If the data was backed up using *Dump* then the filesystem can usually remain mounted in so long as the files to be overwritten are not in use. *Tar* and *Cpio*-based backups can generally be restored in so long as the files they overwrite are also not in use. In certain cases, an alternative to using another boot device for performing a restoration, a system can be brought into single-user mode and restored from there. However, if the data resides over the network, then procedures should exist on how to access the network and data from within single-user mode. In certain cases, where specific files must be restored, but cannot be overwritten while the system is active, a rescue or Live CD or other operating system can be will have to be used.

3.4.3.8 Operating system restoration

Restoration of an operating system's filesystem(s) will generally require booting from an alternate operating system such as those found on a rescue or Live CD or other operating system disk, so that operating system-specific files can be overwritten without crashing the system or resulting in data loss. Replacing existing system binaries and libraries while the system is active could result in system instability or inconsistencies. Other data such as configuration files can usually be restored while the system is active, but only in so long as they are not actively opened for modification. Restoring device files will depend on whether the device(s) are currently in use. If so, then the devices must be disabled or the system must be booted from an alternate device and then restored from media. It is very rare to ever have to restore device files from media as in most cases the Linux operating system easily recreates them automatically at boot-up (using the same mechanisms responsible for operating system hardware reconfigurations).

In general, most operating system restorations can be done without ever having to reboot the system using an alternate device or brought into single-user mode as most of the operating system's files can be replaced without ever affecting the system. The only time this is necessary is when making changes to key system libraries, kernel files and devices where system inconsistencies will likely cause a kernel panic or a total system crash thereby rendering the operating system too unstable for use. In many cases, operating system backups can be restored to an alternate directory on the system and when applications and services can be disabled, the files can be replaced at a another time. This is the most suggested for me for performing a restoration where the system's files cannot be immediately overwritten. Once all but the system's locked files have been overwritten or replaced, the system can be briefly rebooted and the locked files moved to their destination from the location of the alternate directory.

3.4.3.9 Users and applications

User-based data can generally be restored at any time. The only time this may be problematic is when a currently opened file is undergoing modification by the user and/or application or service. In such cases, the file cannot be replaced without causing data loss to the file. It is therefore normal for the system administrator to request that affected users logoff the system by a given time so that the restoration of the affected data can be continued. In cases where the user cannot logoff from the system, the data can always be restored to an alternate directory and then copied over to the appropriate location at another time. If necessary, the system administrator can forcibly logoff a user and shut down his applications and open files. However, this should be left as a last resort to users in a state of non-compliance.

3.4.3.10 Databases

Database applications should always be shut down before restoring any database. Unlike many other applications, database applications should not be left operational even if the databases are currently unmounted. Even unmounted databases on quiescent systems can result in data loss or corruption if a database file is restored while the application or service is running. Database files do not necessarily need to be mounted in order to be considered active by its controlling application or service. However, this will vary by database application or service and as such, it is important to be somewhat familiar with the affected database in order to better understand the potential ramifications of a database restoration.

3.4.3.11 File attributes

Not all backup tools are capable of storing extended file attributes. Depending on the tool used and the data backed up, the backup archive may not have saved the files' attributes. Furthermore, during restoration, it is important to use an appropriate restoration tool that is able to extract and restore the various extended file attributes. The most common type of extended attribute is the ACL and if incorrectly restored or backed up, this could result in the administrator having to manually recreate the missing ACL's, a very time consuming process. *Tar* and *Cpio*-based backups do not store or restore extended file and filesystem attributes. *Dump* and *Restore* save and restore all filesystem attributes.

3.4.3.12 Security, compression and networking

If security is a concern, then it is important that the same security measures be taken during restoration as was done for the backup. If certain applications such as *SSH* or *GPG* were used during the backup, then they will also have to be used during the restoration in order to secure the network and/or data stream. Similarly, if the tools *SSH*, *RSH* or *Netcat* were used for the backup, then they too should be used for the restoration. If encryption was used, then the passwords used must be available for data decryption. Furthermore, if administrative accounts are used for the backup, then the same administrative account will be required for a successful restoration. Access to system directories and direct disk access via raw devices also requires administrative access. Just as with data backups, certain permissions will be required in order to write to various parts of the system and to overwrite existing data. In addition, various security mechanisms may be required in order to perform an "over the network" restoration. These types of restoration

should generally be restricted to administrative accounts. The backup media should not be accessible to just anyone. Only authorized personnel should be able to use and access backup media. Finally, if compression was used during the backup, then the same tools will also be needed in order to decompress the network or data stream.

3.4.3.13 Multi-volume restoration

Multi-volume backups will require multi-volume restorations. It is therefore important that any backup script(s) used also be capable of handling multi-volume restorations otherwise, user intervention will be required in order to perform these restorations. It is also a good idea to place these scripts onto any alternate boot devices so that they are readily available.

Multi-volume backups carried out using disk or tape based software such as *Dump* and *Bacula* can readily handle restoring from multiple volumes.

3.5 Miscellaneous

3.5.1 Errors

Errors can and will always occur, sometimes predictably and other times, randomly. Some errors are caused by human intervention (or lack of it) and others by software or hardware errors. An exhaustive list cannot be provided here since there are too many potential sources of errors to list. Nevertheless, both the backup and restoration process should always be performed by someone knowledgeable of the computer system, the operating system, the applications and services used, the users and the network lay out. An individual such as the system administrator will be well-suited to perform these important tasks and functions.

In addition, the tools used for backing up and restoring data should be error tolerant. Unfortunately, most backup software can only handle the most basic of errors; therefore, human intervention may be required at some point when performing complex backups or restorations. Only very complex and costly backup tools can handle complex errors semi-automatically. However, no single tool could ever deal with all the potential problems that could be encountered.

Errors can also occur in the backup media including both tape and disk. Media errors can be mitigated by performing periodic media tests. The tools examined thus far can generally handle disk-based bad blocks. When dealing with tape errors, *DD* can always be used to forcibly extract data from the backup media. However, portions of missing data from the archive(s) due to bad blocks may result in a partial or complete loss of data. Different tools will handle damaged archives differently. However, in general, all the tools examined are all able to deal with damaged archives and skip on to the next valid portion of the archive. Finally, *Dump* is very tolerant of archive errors; tools such as *Tar* and *Cpio*, less so.

3.5.2 Testing

Before actually proceeding with any backup or restoration, it is important that procedures be adequately tested ahead of time in order to adequately ascertain the safeguarding of data and

systems. In addition, the backup and restoration procedures should be tested and verified in order to determine that they conform to the organization's operational environments and policies. The procedures should also be periodically updated to reflect changes made to the underlying architecture or the introduction of new technologies therein.

Testing should always be done in a controlled laboratory or environment where changes can be made to the systems and network without affecting the operational network. However, once the tests and procedures have been conducted and are considered ready for use, they should then be briefly tested on an operational system. The operational network test should consist of a short series of trial runs to determine if several of the most likely to be used procedures will work under varying conditions. However, this can only be done if the operational environment and organizational policy permit this and if the affected systems can be taken offline. Otherwise, laboratory-based testing will suffice. Once most or all of the tests are conclusive, then the procedures can be rolled out into the operational environment.

3.6 Summary

As examined, performing data backups and restorations can be both a cumbersome and time-consuming process. There are many issues and sources of potential complication that must be analyzed and contended with before actually implementing and proceeding with any specific backup or restoration. After having read this section, the reader should be more aware of the potential impact backup and restoration methodologies will have. However, any methodology developed must be specific to the organization's requirements, policies, computer hardware and operating systems.

In attempting to determine the necessary backup and restoration methodologies, it is possible that much trial and error will be required in order to determine the best overall approach. Furthermore, until definitive backup and restoration policies and methodologies are in place, it may make more sense to use more than one type of backup tool and media type. In addition, the type of backup tool, the media to be used, error handling as well as the type of data and its size will largely determine the overall methodology and approach that will be used and implemented.

4 Methodology II – system maintenance steps and procedures

4.1 Objective

The purpose of this section is to examine the various issues surrounding operating system maintenance, whether it is for updates, upgrades or the manual system maintenance of an operating system. There are many reasons why operating system maintenance should be performed. For example, to fix bugs, improve performance, add new functionality, improve system services, increase system stability or to prepare a system for an operating system hardware reconfiguration (or hardware migration). Fundamentally, this report's objective is to provide the necessary information required to adequately maintain computer systems so that they can be updated, either for its own sake or so they can undergo a hardware reconfiguration or migration (see reports [2, 3]) in order to accommodate for changes in hardware.

4.2 Introduction

Before any system maintenance-related action can be undertaken, many issues will require a thorough examination by both the system administrator and support personnel who oversee the various C2 systems. Both this report and section take a more global overview of system maintenance as compared to Section 3, where precise actions and commands were examined in-depth. Throughout this section, a highly technical discussion is not necessary, as the proposed concepts are of a higher level and. This in turn requires that the reader implement a lower-level perspective in order to apply the proposed methodology to various environments and organizational policies. In addition, it was determined that providing detailed commands and actions for this section could lead to confusion. Linux distributions differ by varying amounts, thus attempting to write a lower-level system maintenance methodology that encompassed them all would be too complex to write in a technical memorandum. However, there are enough similarities among them that a higher-level perspective could be examined.

Periodic and regular system maintenance is necessary to ensure that a computer operating system is both easier to maintain and administrate over the long-term. In addition, a well-maintained operating system will better adapt to periodic hardware changes, facilitating reconfigurations and migrations. Furthermore, regular system maintenance, whether through regular updates, upgrades, code patching or manual recompilation/reinstallation, all form a part of the routine system maintenance necessary for the long-term functioning of any operating system. The type of system maintenance that the reader will employ will largely depend on both the age and type of operating system in use.

This section provides a series of actionable overviews that can be taken up by system administrators and users alike in order to develop a coherent update or upgrade methodology. Each specific organization is unique such that no simple cut-and-paste methodology can be applied at all times. Thus, the work examined herein focuses more on what should be done rather than how it is done. However, because of the approach taken throughout this section, a higher level of UNIX expertise is required in order to perform many of the necessary lower-level system

maintenance-related tasks. In addition, certain low-level tools (see Section 2.4.5.3) will be required for carrying out many of the system administration maintenance-related tasks.

Manual system maintenance, in contrast to updates and upgrades, should only be done on a case-by-case basis, and only when necessary. However, the level of expertise required for a thorough discussion of this topic is far beyond that found herein. By using the aforementioned low-level tools (see Section 2.4.5.3), many problematic issues that occur during operating system updates, upgrades and manual system maintenance, can be resolved by tracking down the root cause of the problem (e.g. library inconsistency, incompatibility, etc.).

However, it is important to understand that at one point, updates and upgrades may no longer be available (for a variety of reasons), leaving manual system maintenance as the only method for maintaining an operating system. Although this is not an easy task, the discussions examined herein will be of immense value when this scenario occurs.

4.3 System maintenance

4.3.1 Reasons for performing system maintenance

There are a variety of valid reasons why system maintenance should be performed regularly (or periodically), regardless of the maintenance type. The following is non-exhaustive list, but it does cover many of the issues important to this section. Consider the following:

a. Reduce the time required for future maintenance.

b. Decrease the system's susceptibility to attacks and exploitation through known vulnerabilities and/or bugs.

c. Improve management of both physical and virtual memory.

d. Provide increased system stability via newer device drivers and provide enhanced hardware features and/or capabilities.

e. Enable a successful operating system hardware reconfiguration or hardware migration.

4.3.2 When and why to perform system maintenance

Before proceeding with system maintenance, in any of its various forms, it is important to determine why and if various operating system components should or need to be maintained. For example, if only one application or service requires maintenance, then it may be more appropriate to manually modify the service or application instead of implementing an update or upgrade that could affect the entire system. Often, deciding whether to carry out a specific course of action will require an impact assessment to determine what will change and whether the changes are justified vis-à-vis the number of changes to be implemented.

Often times, for small tasks, manual system maintenance may be the preferable course of action. Certain software components can be easily maintained manually, while others should be done

using updates and upgrades. For example, kernel recompilation using newer source code is not a difficult process. However, all too often, required kernel features are excluded from the compilation and this can potentially cripple a system. As such, an update or upgrade may be more appropriate, as the kernel has already been configured for specific use with the current operating system. On the other hand, simple changes such as downloading and installing the latest office suite should not cause a discernible impact on the system as a whole (other than new suite features and possible user-based configuration changes). However, each case is different. Three important factors to consider are: (1) the modifications required, (2) their justification and (3) if maintenance is to be carried out, by which means.

There are many reasons both in favour and against any particular course of action. It is therefore highly important to consider both the advantages and disadvantages before proceeding with any system maintenance, even before laboratory testing. Doing otherwise could prove to be a waste of valuable time and resources that could otherwise have been better spent working on and solving other more pressing system administration-related issues.

At certain time, system maintenance should be performed, at others it should not. The following is a non-exhaustive list, but it does cover many of the various issues that are important to this section. Consider the following:

a. How difficult is the maintenance to implement? If the maintenance fix is complex or excessively long to implement, then it should be done using a more automated method such as an update or upgrade.

b. Consider what must be changed. If manual installation of the maintenance fix will cause too many potential changes, incompatibilities or inconsistencies, then these changes should be implemented using an automated approach such as an update or upgrade.

c. What requires maintenance? Does the operating system in general require maintenance or specific components such as the kernel (and its subsystems) or system services and applications? Updates are often more effective at providing newer kernels and system components as they are already precompiled for the current platform and distribution.

d. Is maintenance to be performed against one or more applications and/or libraries? If this is the case, then it is often easier to perform manual system maintenance in so long as the source code does not need to be modified in any major way.

e. Does the maintenance require source code modification? If too much source code is to be modified, then time might be saved by implementing an update or upgrade instead. On the other hand, if there is only a small amount of source code to modify, then manual system maintenance may be the correct choice of action.

f. Is the maintenance necessary to fix one or more specific system bugs? System bugs such as library or kernel bugs should generally be fixed using updates or upgrades rather than manually modifying kernel source code, in order to not introduce new incompatibilities or inconsistencies. However, this will depend on the type of bugs and their scope.

g. If the bugs are application-related, do they cause usability issues? It is important to consider if it is worth updating or upgrading an application simply to fix a couple of bugs that cause only minor inconveniences. Sometimes, newer applications cause more trouble than they are worth (e.g. user retraining). It is therefore important to determine the severity of the bugs before implementing any changes.

h. Will performing maintenance provide any new productivity improvements (e.g. improved functionality or performance or both) in applications and services, or will it improve the overall operating system? What type of changes will be made? Are they far reaching or limited in their scope? Answering these questions will help to determine if the changes should be carried out.

i. Is maintenance required to fix various security-related concerns? Depending on the security issue at hand, it may not be necessary to correct, as it may only be applicable to unused system services or applications. On the other hand, not fixing highly used services or applications could increase the system's overall vulnerability.

j. Are new vulnerability and/or other security maintenance fixes available? If they can be implemented without excessively changing or affecting the system, then they can be implemented using manual system maintenance. If, on the other hand, the maintenance fixes are complex to implement or make too many changes (known and unknown) to the system, then it may be more appropriate to use an automated method such as an update or upgrade.

k. Are the security-related maintenance fixes kernel-related? If so, then it is often best to directly apply the fix to the source code and recompile the kernel. This can also apply to binary patches. However, if the level of changes made to the source code is extensive, then it may be more appropriate to use an automated method such as an update or upgrade. In addition, kernel recompilation and feature selection can be complex, resulting in a crippled kernel and thus, an unusable system.

l. What are the overall advantages and disadvantages of performing system maintenance? If the advantages outweigh the disadvantages, then the system should be maintained using the appropriate maintenance type. However, if there are too many disadvantages using one maintenance type, then another form should be used in its place.

m. It is important to consider when the last maintenance-related changes were made. If the system is relatively up to date, then it may not necessarily be appropriate to make changes to the system for non-critical maintenance fixes. Sometimes, making maintenance-related changes may be more work than is necessary for the level of maintenance to be provided.

n. Are the maintenance changes critical? If so, then they should be implemented unless there is a good reason not to (e.g. services and/or applications will be rendered non-functional). However, if maintenance-related changes are not implemented immediately, then when a more appropriate set is available such as an update or upgrade and it does not cause excessive adverse effects, then it should be implemented.

4.3.3 Requirements for system maintenance

Before proceeding with any system maintenance, there are certain requirements that should be met. No system changes should be made until these requirements are ascertained. Doing otherwise may result in an inoperable system that would have to be restored from backup media. The following is a non-exhaustive list, but it covers many of the issues important to this section. Consider the following:

a. What requires maintenance? Does the operating system require it or does it concern various components, applications, libraries or other packages? Prepare a list detailing specifically what requires maintenance. If enough components require maintenance, then it is likely a good decision to proceed with system maintenance.

b. Determine which type of maintenance is to occur. If only a small number of changes or modifications must be made to the system, then manual system maintenance may be more appropriate. If many changes are to occur, then an update should be used. If updates are no longer available or certain new technical innovations or fixes are available in the upgrade, then it may be better to perform an upgrade.

c. Are all license issues resolved? Has the operating system been re-licensed (this has the potential to happen), as updating or upgrading to a newer version may trigger this. This would cause the modified operating system to fall under the new license.

d. If manual system maintenance is to be applied, then it is important to determine if any changes have been made to the software licenses and their corresponding software packages. Software packages modified by the maintainer's updates or upgrades are the maintainer's legal responsibility. However, software packages that are manually maintained whose underlying license has changed remains the legal problem of the end-user.

e. Is the test laboratory environment ready for use? Is the test computer's hardware and software installed? Are they functioning correctly? Before changes are ever deployed to an operational system, they should always be tested and verified on a non-operational system that can be experimented on.

f. Have diagnostics been run on the hardware? Diagnostics should be run prior to performing an update or upgrade as a hardware bug could cause the entire process to come to a halt and even corrupt the operating system. In addition, certain hardware configuration checks should be made in order to ascertain that the system is in good working order including verifying the system's log files for any hardware-related issues.

g. It is important to verify the system's logs and determine that there are no outstanding errors or bugs unless the purpose of maintenance is to fix those specific issues.

h. Has the software and operating systems on the test hardware been verified to determine if they are functional? Is the operating system in a functional state? A system in proper working condition is easier to maintain than an unstable system.

i. Have backups been successfully completed? It is important that backups be available as any failed attempt at performing system maintenance could require the reloading of system data from backup media. It is important to ensure that backups are done periodically for both the laboratory testing facility and the operational network.

j. Are the test users ready for testing the system after system maintenance is performed? If system maintenance will have a visible impact on the users or the way they work, it will be important to test what changes they will experience and whether the changes are more conducive or a hindrance. If the changes are more of a hindrance, then the changes should be rolled back.

k. Is the system administrator(s) ready to proceed with the system maintenance? Does he have everything required? It is important that the system administrator and his support staff be ready to continue with the system maintenance modifications and that all materials (e.g. update or upgrade media, backups, etc.) are available should they be required.

l. Is the operating system ready to receive an update, upgrade or manual system maintenance? Are applications and services in use disabled? Have users logged off from the system? It is a common practice that before carrying out system maintenance that any potentially affected application or service be disabled and that affected users be logged off. In addition, it may be necessary to bring the system down to single-user mode or even boot from another media if performing an update or upgrade.

4.4 Maintenance types

Although this section emphasizes long-term system maintenance of Linux-based computer operating systems, other types of system maintenance are equally suitable, depending on when the original C2 system was deployed. In general, while the various forms of system maintenance achieve the same overall objective, they are accomplished through different means (e.g. updates, upgrades, code patching, etc.). These various forms of maintenance can be broken down into three general categories: short-term, medium-term and long-term system maintenance. These three forms of system maintenance can be accomplished using any or all of the various maintenance options provided in [2], although there is likely to be some overlap.

These three maintenance types should help satisfy the Royal Navy's key requirement, that a computer operating system be maintainable for the duration of the C2 system's lifetime. This lifetime is likely to be a period of at least 15 years and possibly as long as 25 years. Currently, it is still uncertain which computer operating system will replace the current C2 operating system aboard the Halifax-class frigates. However, when they are refitted, it is very likely that a Linux-based operating system will power the new systems.

4.4.1 Short-term

Short-term maintenance is generally the easiest and simplest type to perform on a given Linux operating system. The operating system is generally maintained using updates provided by the distribution's maintainer. Commercially supported Linux operating systems can easily expect

updates to be available for at least the first 2 to 3 years of the system's lifetime. However, different support contracts with a given vendor may allow for extended periods of operating system updates for possibly as long as 5 years. Updates are normally provided for the duration of the support period in so long as the distribution's maintainer has not switched to a newer version and has ceased support for older versions.

Most commercial distributions are able to directly download and install updates, as they become available, using sophisticated GUI-based applications. Support contracts may also stipulate that rather than use a network-based path to obtain the updates, they be provided on physical media, such as quarterly update CD's.

Updates do not generally cause large operating system disturbance, although depending on the specific update, this does not always hold true. Instead, they usually provide bug fixes and software patches, but can also provide improved application stability, system reliability and performance, and less commonly improve application functionality. Generally, most updates are transparent to the end users, although from time to time, it can occur that certain applications may change slightly from update to another. Normally, this should pose no problem to end users or to system configuration files. In addition, currently supported applications and libraries should not be affected by maintainer-based updates, as they should already have been thoroughly tested by the maintainer. Rarely do updates upset the balance of system libraries, although when it does happen, manual system maintenance may be required to rectify certain inevitable problems. These problems can often be fixed by creating new configuration environments[12] for the affected applications and libraries.

Updates periodically include newer kernels, although drastic changes in kernels (e.g. switching from kernel 2.4.x to 2.6.x) are indeed a rare occurrence. Minor kernel changes do not normally result in changes to system calls and have therefore little tangible impact on the system's existing applications and libraries. Nevertheless, keeping the operating system's kernel up to date will improve a system's ability to undergo a hardware reconfiguration or migration.

Finally, manual system maintenance can be used at any time to manually update or upgrade an application, library, service or kernel. However, due to the amount of additional work required when performing manual system maintenance related tasks its use is likely to be rather limited, although it is always available if it serves as a more viable solution to an update. Manual system maintenance cannot replace operating system upgrades. The cost and effort in attempting this manually may be staggering.

4.4.2 Medium-term

As operating system update-based support is dropped in favour of supporting more recent versions of the operating system, an upgrade will eventually be required to move up to the next level of technology, unless for some reason, manual system maintenance is preferred. Upgrades,

[12] A configuration environment could be one of several things. Normally, it is when an application and its libraries are moved to another location on the system and a new system and/or user configuration is created for it and is independent of the rest of the system. Another type is creating a *chroot*-based environment where the application runs inside of another more restricted environment. Unfortunately, these environments are complex to setup and configure and may not always be an appropriate solution.

however, are usually preferable to manually maintaining and supporting an operating system, due to the complexity involved. Thus, an upgrade is an operating system medium-term maintenance solution. It can be expected that a commercially deployed operating system will likely undergo an upgrade every 3 to 5 years, unless extended operating system update support is stipulated in the support contract with the vendor.

However, simply because update support has stopped may not be enough to justify upgrading, especially if large changes are entailed. Conversely, upgrading may be justifiable due to new technical innovations that provide features including a substantially newer kernel, applications, services and libraries. An older operating system can normally be upgraded with only a few problems to be expected along the way. It is, however, probable that at least several important applications may no longer work as expected, due to application and/or library-based changes. If this occurs, part of the upgrade can normally be rolled back and the missing files recovered from backup. Then a new configuration environment can be created for the affected applications and libraries. However, this may not always fix the problem and sometimes, the only way to fix it is to track down its root cause(s) using the tools detailed in Section 2.4.5.3. In the worst cases, applications and/or libraries that do not work properly can be recompiled and/or modified from source code, but this can be time consuming and complex to the uninitiated.

Therefore, unless there are valid reasons for not upgrading to a newer operating system version when update support has been terminated, in-house laboratory testing should be conducted to confirm that upgrading would result in a stable operating system. In addition, in-house testing will confirm if the changes caused by an upgrade conform to organizational requirements and policy (e.g. changes in operating system security policies).

Once the system and its applications have been deemed functional, the system can then be updated using vendor provided upgrades. Once an upgrade is successful, subsequent updates can then be applied to the system until the next major upgrade is available or necessary. It is likely that updates for commercially supported operating systems will be available for as long as 3 to 5 years before the next major upgrade is required.

Therefore, upgrading is a medium to long-term system maintenance solution. Its caveat is that upgrades eventually must themselves be superseded by a series of updates, only to be followed by another upgrade. With time, through the successive implementation of updates and upgrades, the operating system may begin to show symptoms of software decay. Software decay occurs when enough small changes have been made to a computer system where it no longer behaves as expected due to the number of progressive modifications that have substantially altered the system, its usability and reliability. This is generally caused by an excess of library incompatibilities and inconsistencies that can no longer be easily managed through alternate configuration environments. Although rare, this can sometimes be caused by system call changes in the kernel. This type of issue is almost never seen for updates designed for the original operating system but can occur as time progresses and subsequent upgrades and updates have been applied to the system. This may take as long as 5 to 15 or more years before coming into full effect.

4.4.3 Long-term

The final type of maintenance to examine is manual system maintenance. This form of maintenance is generally suitable for either long-term maintenance (e.g. may occur after 10 to 15 years of continuous maintenance) or when one or more of several possible events has occurred. The most likely events to occur are:

1) The distribution's maintainer has gone out of business.

2) The maintainer has been bought out by another company that either does not support or sell Linux-based products.

3) The maintainer no longer supports Linux as it has changed its line of business.

4) Alternatively, if an organization can no longer benefit from the use of successive upgrades and updates, as they have rendered key applications and libraries non-functional.

Manual system maintenance may be required as a permanent solution after having applied many upgrades and updates to a system. The time required to fix the problems caused by them is equal to or greater than the amount of time necessary to manually maintain the system.

Manual system maintenance can be used at any time during an operating system's lifecycle in order to continue supporting and maintaining it, as well as its applications and libraries. However, this is not advisable, as it should only be used when necessary. It is almost never used after the initial deployment of C2 system as updates can generally fulfill this role. It is only after many years and system changes have occurred that it should be seriously considered. However, this is not to say that every software package must be manually recompiled and reinstalled. Manual system maintenance can be carried out using the manual application of updates and upgrades without using any GUI installation program. As such, the system administrator becomes responsible for ensuring that required software packages are up to date through whichever means are at his disposal. Equally applicable is the implementation of new or modified source code, its recompilation and reinstallation. In so doing, the system administrator gains additional control over which applications and libraries are changed.

What sets long-term manual system maintenance apart from the other types is that the system maintenance is now almost entirely dependent on the system administrator, the method he chooses (as the situation requires) and his skills and expertise in maintaining the system. In addition, it is very likely that the system administrator will not be able to carry out all the required tasks by himself and will have to work with other technical support staff. However, in so long as the applications and libraries in use continue to be supported by the vendor, manual system maintenance may only be required occasionally. Fortunately, it is only necessary when applications and/or libraries are no longer functional, due to various sources of conflict.

Manual maintenance is unfortunately the longest and most complex form of system maintenance possible and generally requires extensive in-house testing. Sometimes, it will require the direct modification of existing source code and at other times, the application of patches or simple

downloads and installations. Cases will vary widely and many situations are likely to be unique and equally cumbersome and complex to implement.

Long-term manual system maintenance is not ordinarily required, so long as other maintenance options remain open and available [2]. It is entirely possible that through highly specific support contracts, customized updates and upgrades can be provided from the vendor that will be able to adequately satisfy the Royal Navy's long-term requirements. However, things tend to change and without at least giving this option due consideration, the Royal Navy may find itself without any other viable option other than to deploy a completely new operating system from scratch instead of building upon and reusing available resources.

4.5 Licensing

Software licensing can be a contentious issue and it can have a dramatic impact on which type of long-term maintenance to use. Almost all FOSS-based software is distributed with a specific type of license. There are many different types of open source licenses and although many of them are similar, each one grants specific rights and limitations to the end-user. Commonly, the end-user is considered the organization that is using or deploying the software in question. However, this definition applies equally to individuals using the same software at home.

It is common for software licensing to become a complicated and convoluted matter not just for government departments, but also for any involved party. The problem with software licenses is that the language they employ is often subtle, vague and too often all encompassing. Unfortunately, where software licenses are concerned, there are often many subtle nuances that must be understood in order for the end-user to understand his rights and obligations. In general, licenses tend to grant few rights although they do typically contain many limitations. Report [3] may serve as an interesting starting point for those interested in comparing two different commercial FOSS-based licenses. In report [3], a table has been provided to serve as a useful comparison.

Nevertheless, FOSS licenses tend to be less restrictive than their commercial counterparts are. Commercial Linux distributions, while open source in nature, are typically bundled with commercial licenses that are often more restrictive than most FOSS-based licenses. However, these licenses tend to not be as confining as other non-FOSS commercial licenses.

4.5.1 Types

An important question to ask before deploying just any open source software is what type of license is the software bundled under? There are currently more than 50 different types of open source licenses and while many are similar, they each have their own specific advantages and drawbacks, rights and obligations. The two most commonly encountered are GPL and BSD-based. Many of the other currently available licenses are derivations of these two types.

In a military setting where intellectual property (IP) is important, it is advised that a BSD-based license be used as it allows the end-user (e.g. the Royal Navy) to preserve full rights to any changes made to the source code. Conversely, GPL-based licenses generally require that all

changes made to the source code be given back to the community. Where matters of national security are concerned, this may not be in the best interest of Canada.

It is therefore important to understand the implications of license types and their impact on C2 system maintenance. If maintenance is provided by the vendor [2], then all legal responsibility rests therein. Otherwise, if maintenance is done in-house, then a legal analysis of a given license should be considered before adopting it. Unfortunately, the Canadian government has not yet reached a decision about how to treat open source licenses (this is normally the jurisdiction of Justice Canada).

4.5.2 Compatibility

License compatibility issues are important but are often overlooked, for both FOSS and proprietary software. An important question to consider is 'does the software license allows for making modifications to the operating system, application or service?' Moreover, as software changes over time, so do the underlying licenses. It is common for FOSS software to change to another license type. This gives rise to the following question: 'how will changes in FOSS licenses affect subsequent modifications of that software?' License changes occur in large part due to pressure from the software's users and developers, but sometimes it occurs if the software maintainer has a philosophical "change of heart." While this is very unlikely to occur for well-established FOSS projects such as the Linux kernel, for core system tools and utilities, the possibility always exists. License changes may also be necessary to ward off claims of intellectual property [19].

If system maintenance is to be required in the near future, it is important to determine if current software licenses are the same or compatible with their newer counterparts. The answer will depend on many factors, the most important of which is the governing license that all other software licenses fall under. For example, a commercially based distribution may have a governing license stating that its license supersedes all other underlying licenses. Thus, all changes to the software and underlying licenses are the responsibility of the vendor, not the end user. Therefore, if a software package changes its license to another type that is incompatible with the current governing license, it remains the vendor's problem.

Conversely, those performing manual system maintenance and no longer have a valid support contract with the vendor fall outside the protective umbrella of the vendor's governing software license. Thus, should the underlying software licenses change from one type to another, then legal counsel should be sought out before proceeding to make any changes. Generally, changes to FOSS licenses, although minor in nature, will continue to allow the end user to work with and modify the underlying source code. However, in the event that a new license is incompatible with the needs of the end user (e.g. divulging intellectual property) he can, at his discretion, choose to replace the software package for another similar type that satisfies his requirements. The other alternative is to not adopt the newer software package with its incompatible license and manually maintain the older software, without using any source code from the newer software or from the user community. For example, consider a FOSS program originally developed under a BSD license that is eventually re-licensed under the GPL license. This would significantly affect new modifications to the source code and require them to be distributed back to the community.

4.5.3 Permissions and limitations

It is important to understand what is permitted under both a governing commercial software license and the various underlying FOSS licenses. Each license grants a specific set of rights and obligations. Some licenses, commercial and open source, are very restrictive while others are far more permissive. For example, GPL-based licenses are far more restrictive than their BSD counterparts are. However, much will depend on what is specifically required from the license and software, including any potential changes to be made to the source code, both in the present and future. For example, if a software component's license has changed and no longer reflects the organization's needs, it may be more appropriate to replace it with another component whose license is more in line with the organization's requirements. Thus, the rights and obligations sought after in a license should closely reflect the organization's requirements from the software and C2 system. It is also important to consider that a governing license may change over time as per implemented updates and upgrades. Therefore, accepting newer versions of software can have a direct impact on the type of maintenance to use, which can affect the C2 system's long-term evolution.

4.5.4 For consideration

As with all legal issues, operating system and software licenses vary considerably, from distribution to distribution. The vendor may simply apply a standard generic license to the distribution, thus covering all the underlying software with its license. Some of the issues that have not yet been examined but that are important to consider can include the following:

a. If a distribution is bundled using a standard generic license, is manual system maintenance permissible? If it is permissible, what are its limitations? If not, then legal counsel should be sought out before attempting manual system maintenance.

b. Depending on a vendor's license, it may be necessary to renew on a yearly basis the software support contract simply to have the right to use the distribution, let alone modify it [3].

c. Do the open source licenses found with the distribution allow for the modification of the kernel, system configurations, applications and libraries? Generally, this is not an issue. However, this may be problematic if source code and IP are integrated therein.

d. Have any of the software packages' licenses changed? If they have, are they compatible with existing licenses and with both current and future modifications to source code? What is the impact of these license changes?

e. Can a vendor's distribution source code be patched or integrated with non-vendor source code that has been provided by the open source community or third party?

f. According to the governing license, can specific packages be changed, modified or replaced? What do the licenses of the individual packages permit?

4.6 Laboratory testing

Before considering the implementation of a system maintenance-related action on an operational C2 computer system, whether it is an update, upgrade or manual system maintenance, all actions should first be thoroughly tested. Specifically, tests should be conducted within the confines of a laboratory environment where tests can be safely performed and evaluated without jeopardizing the stability of the operational network. A laboratory is an exceptional setting for testing patches, bug fixes, updates, upgrades, source code modifications, etc., before they are ever rolled out into an operational environment. Laboratory testing provides an opportune time and location to perform system reaccreditation and recertification in order to minimize the overall impact that would otherwise be experienced in operational environments. Since the Royal Navy has expressly stated that all changes must be certified and tested, this setting provides a unique opportunity for carrying out these necessities. Once a test system has been reaccredited and recertified in a laboratory setting, it can then be safely implemented (barring certain precautionary safeguards) onto operational systems.

In addition, those interested in software degradation and operating system evolution will appreciate the use of such facilities to study and examine how operating systems and their software respond to many such maintenance changes progressively made over the years. Others may some may see laboratory testing as a waste of time, especially in testing small and seemingly innocuous changes. Nevertheless, even small changes made progressively overtime build up and can cause a "snowball effect" which can effectively cause seemingly functional systems to cease functioning.

4.6.1 Laboratory

Before proceeding with any laboratory tests, it is important that the testing environment be as similar as possible to the operational environment. The laboratory does not need to have the same number of computers or users. However, the laboratory should utilize the same type of telecommunication equipment used in the operational network such as radio, routers, bridges, uplinks, etc. The physical computer systems should be of the same make and model as those in the operational network, as software can respond differently when used on different systems. This is in fact a well-known issue of computer systems, particularly hardware dependent software such as kernels and device drivers. In addition, when troubleshooting, this will help to isolate software related issues caused by incompatibilities and inconsistencies in hardware. Furthermore, the operating systems, applications, services, etc., should be similarly configured as per the systems they are meant to represent on the operational network.

The use of similar environments will help to "shake down" test systems and reveal software and hardware bugs before they are encountered in the operational network. Furthermore, a laboratory will provide a more realistic environment for system administrators, support staff and users to test and experiment with the various systems. This will help to determine if applications, databases, files, telecommunication systems and other various services and resources are equally available, responsive and functional as they were before any system modifications were made. In the end, all this testing will help to enable a faster and more simplified deployment and transition of the required changes onto an operational setting.

4.6.2 Laboratory isolation

It is important before proceeding with any tests or modifications that the test environment be completely isolated from the operational networks. Both networks need to be free from contamination that could be caused by the other. If either environment were interconnected, it would make bug tracking and troubleshooting more problematic. Even environments separated by NAT-based firewalls or VLANs cannot stop all possible cross-contamination issues. Thus, when conducting tests of C2 systems and their networks, it is important to determine the sources of possible outside influence in order to track down bugs or other issues. Furthermore, the extra layer of objectivity afforded by isolation will make system reaccreditation and recertification simpler as external sources do not have been taken into account.

4.6.3 Backing up

It is important to back up all data before proceeding with software testing in the event changes must be rolled back. However, depending on the specifics of the test environment, the software used and installed, backups may not always be necessary, although they are often a good idea. In addition to preserving data, backups will allow for the additional testing of new backup technologies, methodologies and emergency restoration procedures.

Data to be backed up will vary considerably and will be according to the type of data and test systems to be backed up. A backup may consist of user data, applications and configurations or it could be a full system backup. However, the backing up of multiple test systems may result in excessive work and it may therefore be more appropriate to use cloning[13] and distribution systems instead. This approach is particularly well suited to environments where all or most of the systems are identical. More information on backing up can be found in Section 3.

4.6.4 Benchmarking

Before deploying any successful changes made to the C2 test systems onto the operational network, it is important to consider the performance-based issues that can inadvertently affect the operational systems. One way to determine if a set of "successful" modifications will cause inadvertent changes on the operational network is to test them in some way. The standard method of testing consists of examining application and service-based functionality as well as usability tests. Another method that is proposed herein is benchmarking.

Benchmarking is a performance-based test that measures the performance of the system, application or service against some measurable unit. The most commonly used unit is time, although any other useful performance-based unit is equally acceptable. Benchmarking is useful because it can help to pinpoint slowdowns caused by new software or set of modifications. It can verify if performance (could be the system as a whole, an application or service) is similar to before the changes were implemented. Through benchmarking, it will often be possible to determine if a series of changes was beneficial or has degraded overall system performance.

[13] These software systems do not generally perform bit-copies. Instead, they usually copy files and associated filesystem metadata. This type of filesystem copying and distribution may be suitable for some environments but not in others.

Modifications can be tested individually or in groupings. Furthermore, it is possible to benchmark software according to changes in both the software itself and its configuration.

By benchmarking different modifications and attributing a performance-based score to pre- and post-modification systems, it is possible to determine with good accuracy whether a given change or set of changes is successful. A general rule of thumb is proposed: if the performance of a system, application, etc., is the same or very similar to the original, then the change(s) can be considered successful. Conversely, noticeable slowdowns can be indicative of problems requiring resolution or an inadequate or incompatible change or set of changes having been made.

The theory behind benchmarking is that non-effectual or counterproductive changes and modifications are more likely to cause system, application and service slowdowns. Thus, if a set of modifications actually causes an unexpected or significant slowdown where none was previously seen, then it is likely that those changes were either incorrectly made or are detrimental to the system and should be reversed.

However, in order for benchmarks to be useful, a base score is required and this can only be accomplished by benchmarking the original system marked for deployment. The system, as well as key applications and services, should be benchmarked and serve as a comparison for future benchmarks. In addition, benchmarks can vary widely due to extraneous factors. Thus, benchmarks should be repeated several times in order to average out the results. In so doing, benchmarking can help to objectively pinpoint potential performance gains and issues that are the result of one or more modifications.

4.6.5 Incremental changes

It is important when testing systems in a laboratory that wherever possible changes be made incrementally. However, this is not always practical or possible to put into practice. When testing various maintainer-provided updates and upgrades, depending on how they are to be implemented, incremental implementation may not be possible, although with tweaking, it often is. Certainly, this is easier to achieve when using manual system maintenance as compared to maintainer-based updates and upgrades, as these generally require an installation program that does not always make it possible to implement incremental changes. However, most distributions provide the ability to fine tune system and application updating. Upgrading, on the other hand, is often far more problematic. Nevertheless, this will vary considerably by distribution to distribution and from version to version. At the same time, configuration files can also be changed incrementally, making a set of changes and testing them before proceeding with another set of changes. These changes can be small or large and consist of one or more configuration files.

Making changes incrementally is always a best practice. However, time does not often permit for this. After each modification (or set of), it is appropriate to perform benchmarking tests and comparisons. Incremental changes can allow several goals to be attained:

1) It enables a more precise targeting and tracking of problems, instabilities and inconsistencies that arise because of changes and/or modifications.

2) Facilitates the rolling back of changes, as there is less to be removed and undone as compared to a full system update or upgrade.

3) It makes version and system change control easier to track and maintain.

Unfortunately, during this process, much will depend on the distribution itself as different operating systems use various approaches to apply updates and upgrades. In general, upgrades tend to make incremental testing difficult and sometimes, for all practical purposes, infeasible.

4.6.6 System administration testing

It is obvious that the system administrator will know the various systems, infrastructures, telecommunication equipment and operating systems very well. It is his job to understand them and to be comfortable with them. That is why, although obvious to state, that the system administrator plays a key role in testing the system after changes have been made to it. The system administrator above all others knows what to expect and how the system in general should behave, including its performance, reliability, security configurations and network-based access and resources. Of course, user testing is also very important. Before user-based testing is done, system administration testing should take precedence. Only after the system administrator finds the system to be functional should other tests be conducted on it. Different system administrators, according to their skills and years of experience, may determine a system's suitability for work using different tools and techniques. While most vendor distributions provide similar UNIX-based tools, they can sometimes behave differently. Thorough documentation of system administrator-based testing is as important as any other test. In addition, results from the various tools, including system metric-based information, should also be included in any documentation.

4.6.7 Behaviour and functionality

Before accepting a set of modifications, several issues are very important to consider. The most important is to determine if behavioural changes are the result of modifications. Behavioural changes could be indicated by a change in screens or system messages when the system boots up or powers down. There could also be various messages written to the console that could be the result of one or more buggy device drivers. Applications and/or services that were once fast and responsive are now slower or unresponsive (e.g. benchmarking). Noticing behavioural changes often is not an easy task, but someone such as the system administrator should be familiar enough with the current hardware and operating systems to recognize many differences. Many system changes can occur after an update or upgrade; this is why, when possible, changes should be made incrementally, tested and observed. Thus, the system administrator is uniquely qualified to determine if aberrant behaviour or functionality is a result of the changes made to the system.

However, the system administrator is not likely to recognize changes made to the various applications used on the system(s) by the various users. Different users will use different applications and services to start and complete their assigned work and tasks. Test users are uniquely capable of determining whether applications and services are behaving and functioning correctly. Test users should consist of advanced users who are fully capable of performing their tasks with the least amount of system administrative support so that they can independently verify if adverse changes are present.

Finally, aberrant system, application or service behaviour is not always caused by modification to the actual application, service executables or libraries. Sometimes, it is caused by changes to the configuration files. Other times, it is caused by a change to one or more library dependencies or system calls.

4.6.8 User-related system changes

User-related system changes must be evaluated in order to determine whether any of the changes made will cause system, application or service-related disruption or failure. Any type of disruption or failure could adversely affect the way users work and interact with both the system and each other. This will in turn adversely affect both their day-to-day tasks and their overall mission objectives, many of which have to be accomplished together as a team. In a mission critical environment, any disruption could be potentially disastrous.

Therefore, it is very important to test the system after a set of changes, even if they are small. Although the system administrator is generally able to distinguish adverse effects to the operating system itself, only the users are uniquely positioned to test the system's applications and services. Of course, not every change is necessarily bad. In fact, most often, changes are necessary so to fix bugs and provide newer or enhanced features.

Laboratory-based testing in an isolated environment will make it easier to determine if users experience any differences in their day-to-day activities and use of the system before changes are eventually deployed onto operational systems. This provides users an opportunity to voice themselves beforehand rather than be forced to accept non-functional or aberrant modifications or changes to their applications and work methodologies. Thus, user-based impact studies are necessary in order to assess the usability of the changes made. Some tests users should consist of "power users," users who are generally self-sufficient. System administrators in general make poor test users, as they often bypass security or organizational procedures to accomplish their task(s) or fail to understand how application-based changes will affect users.

User-based changes should be tested on a case-by-case basis and not all changes require test users. Operating system changes that are sure not to affect applications and the users do not necessarily require user testing. Of course, the system administrator should test these changes.

4.6.9 Impact assessment

For every set of changes made to the system, an impact assessment should be conducted. However, the implementation of an impact assessment should be commensurate with the amount of time necessary to actually perform it, as well as the time required to make the changes. For example, if several very small changes are made and the consequence(s) of these changes are already known ahead of time, then it may not be necessary to carry out an impact assessment. Judgement and common sense should be utilized at all times. Otherwise, the impact assessment portion of testing could become excessively complex and cumbersome.

Before conducting any impact assessment, two questions should be examined. The first is "will the system(s) and network(s) continue to behave as they always have after the changes are made?" The second question applies more specifically to system behaviour and is "if a set of

changes makes no visible changes to the system or the users' ability to interact with it, then is it worth performing an impact assessment?"

Impact assessment-related issues are important to determine. Unfortunately, not all changes and their impacts can be known or understood ahead of time. However, many are known as they are already well documented and may even have been previously tested at a different time or place. Nevertheless, impact assessments should be conducted from within a laboratory-based environment. In addition, this laboratory is critical in order to adequately perform and test patches, updates, upgrades, manual system maintenance and bug fixes in order to determine what potential incompatibilities, incoherencies or instabilities could be introduced into the operational environment or infrastructure. These issues can be caused by changing and/or replacing key operating system libraries and interdependencies, applications, system and application configurations as well as configurations. Furthermore, it is important to determine if any of the changes break or modify system or organizational security policies. Impact assessment testing should be done in conjunction with benchmarking and system administrator and user-based testing.

4.6.10 Modifying system configurations

It is important to consider the impact of system configuration modifications. Some changes are innocuous and are likely to result in no noticeable changes or behaviour to the system. Others may cause services and applications to act differently than before. The modification of system configurations can dramatically affect users and their interaction with the system.

Certain issues should be examined during the different testing phases after an update or upgrade has been implemented. It is important to determine which files have been changed, particularly configuration files. Changed configuration files should be compared to their previous incarnations to determine if any important service or system-wide changes have been made or propagated. This is another reason why it is also important to conduct backups prior to implementing updates or upgrades. Furthermore, this is why documentation, change assessment and versioning control are important.

4.6.11 Outcome testing and manageability

It is important to determine if a set of changes made to the system have resulted in a desirable outcome. However, it is important to define beforehand what should be considered a desirable outcome (e.g. success or failure). Failure could be defined as the causality between a specific set of changes and a disruption in services or applications, system stability or reliability. It could also be defined as an unacceptable change in system behaviour, performance, application or service use. Only after thorough examination and stringent testing, can the cause of a problem or failure, if it occurred, be determined. Then, depending on its manageability, it can be deemed a success or failure. This, generally requires testing on the parts of both the system administrator and test users. Both have a complex job ahead of them. However, it is often more difficult for the users to determine whether their applications and work methodologies will continue to work appropriately. This can include their ability to establish access to and work with data and other repositories, system services and applications, as well as accessing and using remote systems and

devices (e.g. printers, etc .). Ulti mately, if things do not work as they should, it m ust be determined where they fall short.

On the other hand, if the c hanges and the overall i mpact on the sy stem are considered acceptable in that the y have caused little to no discernible issues or disruptions, then the outcom e for the changes c an be considere d a suc cess. The c ase is further solidified if a set of changes i mparts additional be nefits such as i mprovements in usability, perform ance, security , robustness, etc. , without adversely affecting the system.

Outcome testing is not an actual testing phase per se. It is rather a culmination of results obtained from bench marking, im pact analy sis, user tests, system administration relate d tests, etc. All successes and failures should be thor oughly documented as well as justified for an apparent success or fai lure. It is to be expected after some updates and especially foll owing an upgr ade, that there may be some failures. However, in so long as they remain manageable, there is no need to consider the update or upgrade a failure. Ma nageable failu res can still be included i n the update or upgrade process, while those that cannot be reasonably managed should be left out or fixed, if possible, using manual system maintenance.

4.6.12 Versioning and change control

Versioning control is rather sim ple to carryout in s o long as the proper preparations have been made (e.g. versioning control software has been installed). Many software packages are available that can perform versioni ng control for various UNIX and Linux-based sy stems. This software allows for the analy sis and determ ination of which files have been changed, b y whom and when. Some can even go as far as co mparing changes aga inst archived copies. This information is important in order to assess which process (update or upgrade) ha s made which changes and the nature of the change.

However, before versioning control can be implemented, a baseline must be established. Most if not all versioning control software requires a baseline to be establ ished. A baseline can be used to establish information about a system's files such as size, ownership, permissions, creation date, access date, modification date, etc. It i s generally possible to spe cify which files or set s of files should be ta ken into account while c reating the baseline (e.g . specific co nfiguration f iles, directories, binaries and libraries, etc.). With this information, it beco mes possible to determine which files have changed.

Versioning control information is generally stored in a database file, which tends to text-based. The baseline data file should alway s be cons idered as an important starting poi nt for any documentation that is to be written up. In a ddition, thr ough t horough versioning control and documentation, it will be possible to m aintain an established list of known changes that can be used to help track down software and configuration errors as a troubleshooting aid.

4.6.13 Library and kernel modifications

Updates tend to be smalle r in scope with respect to the changes they m ake ag ainst a partic ular system. Upgrades, on the other hand, tend to not only change applications and services but various s ystem libraries and the kernel too. Depending on the t ype of s ystem maintenance

utilized, system changes may be minor or widespread. Generally, library and kernel changes tend to provide additional functionality, improved security, features and bug fixes. Generally, commonly used features such as API's and system calls are rarely removed, although it is always a possibility. When time and resources permit, it is always best to ascertain the specifics of these changes. Using a versioning control system it will be possible to isolate changed files from those left unaffected with relative ease (assuming the files have been baselined). While it is not often necessary to examine these files in-depth, if problems or other issues should arise as a direct result of an update or upgrade, then the changed files can be examined.

Changes to system calls are rather easy to determine, as the kernel's source code is readily available. Libraries, on the other hand, are generally more difficult and time-intensive to analyze, although there are tools directly designed for this purpose. It may be appropriate to analyze libraries only, if as a direct result of their modification, one or more applications or services malfunctions or fails. A full listing of these tools is available in Section 2.4.5.3.

In contrast to applications, services and the kernel, wide scale system-based library changes are considerably more difficult to examine thoroughly, largely because of the far-reaching changes libraries can impart on the system due to intricate and often subtle interdependencies. This is particularly true for system-critical libraries such as the *libc* (C library), which provides most of the operating system's and applications' C calls and functionalities. Discovering which applications and services would be affected by such a change would be cumbersome and time consuming. Therefore, it is important to decide if and under which conditions this analysis is to be conducted.

4.6.14 Reconfiguration and migration

Once all the tests have been conducted and it has been determined that the overall outcome thus far has been successful, then if appropriate, it is time to proceed with an operating system hardware reconfiguration or hardware migration. This step should only be conducted if new hardware has been introduced onto one or more of the test systems. For systems that have not experienced a change in hardware, then this step should be altogether skipped. In certain circumstances where the kernel or libraries have not been changed, a reconfiguration or migration can still be done in so long as the kernel supports the newer hardware or that a third-party device driver is available.

Unfortunately, due to the proliferation of many diverse Linux distributions, there is no generic approach to performing a reconfiguration or migration. Most modern Linux distributions have their own particular method for detecting hardware changes and making the appropriate operating system and configuration file changes. This topic was examined in Report [1].

4.6.15 Documentation

The importance of documentation cannot be overstated. It is important that at least one individual, preferably the system administrator (or other similar person), document information about the various changes and tests carried out. Documentation should be written in a clear, understandable and objective language. The documentation should include, but not be limited to versioning and change control information, as well as changes and other information relevant to

library, application and system interdependencies. It is also important to include listings of changed files and packages, configuration file modifications, as well as kernel and driver changes. Equally important are the various tests that have been performed: impact and outcome testing, benchmarking, behaviour testing, as well as system administrator and user-based tests.

The documentation should convey all necessary information about the changes experienced by a system. The information collected for documentation purposes will be vital to system reaccreditation and recertification, as every system change, modification and overall impact will already have been conducted and detailed. Therefore, if the changes made by an update or upgrade are successful and thoroughly documented, then deployment of the approved changes will be seamless. Documentation will be a requirement in order to gain approval for deploying a set of changes onto operational systems.

Documentation should also consist of problem resolution and other successful troubleshooting techniques that managed to resolve issues caused by an update or upgrade. These problems are likely to occur again when the changes are deployed onto operational systems. Thus, without this information, the process will be more cumbersome and time consuming.

All results, whether good or bad, should be documented. A case can then be made available in the documentation detailing the reasons why an update or upgrade should be allowed to proceed. An objective analysis of all tests, changes and documentation will facilitate the approval process. It is likely that many organizations will have their own procedures and policies for writing technical documentation and approval must have be given before proceeding with any operational deployment.

4.6.16 Approval process

At this point, once all tests have been appropriately conducted according to requirements, time constraints, available testing resources, as well as documented and troubleshooted (if necessary), an overall outcome should be apparent. Regardless if it is positive or negative, a case should be made why a set of changes (e.g. update or upgrade) should or should not be deployed. Using the objectively written documentation, as well as the overall assessment put together by technical personnel, management can make an informed decision about whether or not to proceed with a deployment. The importance and quality of the documentation provided to management cannot be overstated, as its decision will be largely based on the conclusions and findings found within the documentation. Once approval has been given for deployment, a plan should be developed and a course of action put into place for reaccreditation and recertification.

4.7 Deployment

After approval, there is still much planning that remains to be done in order to determine which systems the modifications are to be deployed onto and the order of operations and priority. It will also be important to coordinate deployment efforts with appropriate IT personnel so that they are available for deploying the approved changes as well as for troubleshooting if necessary.

4.7.1 Backup s

Prior to deploying approved changes onto operational systems and networks, full backups should be conducted so that if necessary , sy stem stat es can be rolled back if one or m ore sy stem deployments should go awry . Althou gh the appr oved changes have been thoroughl y tested in a controlled laboratory environment, the possibility still exists that some type of failure could occur during t he d eployment. Such a failure could r esult in network -wide disruptions, potenti ally leading to the unavailability of imported services and capabilities. Therefore, by backing up any system that may be affected by the deployment, it will be possible to restore them to their original state. A backup and restoration methodology can be found in Section 3 of this report.

4.7.2 Deplo yment plan

It is i mportant to develop a deploy ment plan. The plan will examine many issues that must be resolved in o rder to appropriately organize and deploy the various updates and/or upgrades. Not doing so leaves many complex variables to chan ce and can considerably increase the probability that a deployment effort will go awry . A non-exhaustive list of potentiall y contentious issues to consider has been provided below:

a. On how many sy stems will the update, upgrade or manual system maintenance b e deployed? Is the deplo yment to be don e in small or large groups, or done across many systems simultaneously?

b. Are there enough available resources to proceed with the deployment?

c. Does the organization have a policy in pl ace for deploy ments? Does the proposed deployment coincide with existing policy?

d. What will be the effect s on the other sy stems that will not be directly involved in the deployment and will not receive a given set of modifications?

e. Will the operational network or infrastructure be destabilized by the deployment? It may be necessary to tem porarily disconnect the network so as not to affect other sy stems, networks and services.

f. Will the deploy ment be perform ed during tim e allotted for routine sy stem maintenance (weekends, holidays, etc.)?

g. How will deploy ments be done? Will update or upgrade-based deployments be done in part, starting with the kernel, then libraries, services and fi nishing with ap plications t o minimize user downtime and im pact? Alternatively, will the dep loyments be done by implementing the complete update and/or upgrade at the same time?

h. Can users continue to work during the deployment? Will they be able to access their data and applications and perform their routine tasks?

Once the deployment plan has been developed, it should be approved by both management and reaccreditation and recertification officials. The modifications as laid out in the deployment plan are made in the following section.

4.7.3 Rollout

The deployment should proceed according to the deployment plan. It should be developed in terms of resource availability as well as organizational policy and procedure. It is here that the actual modifications to operational systems are made.

Due to extensive laboratory testing, few unknown problems and other issues should be encountered during the operational deployment, although the possibility of this occurring exists. This is because laboratory testing cannot take everything into account that may be found in operational settings. It is conceivable that a deployment could interact in unforeseen ways with the operational network. However, such inevitabilities including electrical[14] surges and shortages, cosmic rays[15] and solar flares, etc., can often be taken into account and largely rectified by using specialized equipment and shielding.

Regardless of the potential for failure (which should be relatively low), the deployment should follow as closely as possible the development plan put together in the previous section. The rollout phase should be documented just as it was for all the previous steps because should failure(s) occur, using well-written documentation it may be possible to track the root cause of the problem or failure.

4.7.4 Reconfiguration and migration

As stated in Section 4.6.14, a reconfiguration or migration is necessary only if hardware changes have been made to one or more of the underlying systems present on the operational network or infrastructure. If no hardware changes have been made, then neither a reconfiguration nor migration is necessary. However, if changes have been made, then in order for that new hardware to function correctly, the operating system will have to recognize it. However, each distribution is unique and each operating system will have its own mechanism for detecting hardware changes and make them available to the operating system. It is beyond the scope of this report to directly examine or detail the specifics concerning reconfiguration and migration. More information can be found in Report [1].

4.7.5 Reaccreditation and recertification

If all the previous steps have been appropriately completed and the changes have been deemed successful and the documentation is objective, accurate, up to date and provides extensive coverage of events, then testing, reaccreditation and recertification should be a rapid process. While the process varies according to organizational policy, its ultimate goal is to determine what has changed and what will be the impact on the system and network. By performing all of the

[14] The use of surge suppressant technology such as UPS-based systems can greatly reduce electrically induced disruptions.
[15] The use of ECC memory can help as can the use of radiation shielding and insolation.

aforementioned steps, these time-consuming verifications are taken out of the loop of IT security personnel and left in the hands of those better able to determine the cause and effect implications of the various system changes that have been made. In addition, once all appropriate evidence has been collected and corroborated from the previous steps, overall, there should be a high degree of certainty about system reliability and stability. The systems should then be deemed satisfactory for operational use and given final approval by reaccreditation and recertification officials.

4.7.6 Wrap-up

Once the deployment has been successfully completed, reaccredited and recertified, it will require a "shakedown" period where hidden bugs not found during laboratory testing or deployment can be worked out. A successful shakedown could take several weeks to several months to complete and learn about any new undocumented features about the implemented changes. During this time, it is important to document any lessons learned (if any) and examine any last minute changes, tweaks or modifications that must be made to accommodate for the modifications and changes.

5 Conclusion

There is no clear-cut or definitive methodology for carrying out system backups or performing system maintenance. The purpose of this technical memorandum has been to propose two methodologies to aid system administrators in these tasks. Different operating systems will require different methodologies. However, because the Royal Navy is interested in deploying Linux-based systems as their new C2 systems aboard the retrofitted Halifax-class frigates, only Linux has been examined. The material presented herein is applicable, in general, to both Linux and UNIX-based operating systems.

The first methodology, found in Section 3, examines the various techniques and technical issues surrounding performing quality system backups before testing and deployment of system maintenance. Different maintenance types and their consequences have been examined in Section 4. Therein, a system maintenance methodology was developed to aid in the testing and deployment of various types of system maintenance.

The Royal Navy has expressed their interest in maintaining the same operating system throughout the lifecycle of their C2 system. As such, they will inevitably have to perform system maintenance and when they do, the issues examined herein will be of great use to both their system administrators and other technical personnel. Many may not agree with the contents herein.

To the author's best knowledge, this is the first document of its kind, as no other system maintenance methodology could be found specifically for Linux or UNIX in general. Although this technical memorandum is slightly similar to document IEEE Std. 14764-2006, the latter refers to software-based life cycle maintenance from a software engineering perspective, while this specific technical memorandum instead examines operating system lifecycle maintenance. Thus, both these are independent of one another.

In conclusion, although Linux distributions vary greatly according to their market niches, they all share certain features and similarities. It is based on these similarities and features that the methodologies in sections 3 and 4 have been proposed. Of course, they are open to interpretation and can be changed to suit different requirements and environments, but at their basis, they offer relevant and useful tips and advice. Section 2 provides useful background information concerning system maintenance, operating systems and the various types of dependencies likely to be encountered.

References

[1] Carbone, Richard. Operating system hardware reconfiguration: A case study for Linux. Technical memorandum. Defence R&D Canada. TM 2006-595. November 2006. http://cradpdf.drdc.gc.ca/PDFS/unc56/p527008.pdf.

[2] Charpentier, Robert and Carbone, Richard. Life-Cycle Support for Information Systems Based on Free and Open Source Software. Revision 1.0. Technical Paper for 11th ICCRTS. Defence R&D Canada. June 2006. http://www.dodccrp.org/11th_ICCRTS/abstracts/136.pdf.

[3] Carbone, Richard. Enterprise Linux licenses: A comparison of licenses between Red Hat and Suse Enterprise Linux. Technical memorandum. Defence R&D Canada. TN 2006-573. October 2006. http://cradpdf.drdc.gc.ca/PDFS/unc53/p526349.pdf.

[4] Wikipedia. System call. Online encyclopaedia. Wikimedia Foundation Inc. October 2006. http://en.wikipedia.org/wiki/System_call.

[5] Carbone, Richard. Does Red Hat 5.0 Support Hardware Refreshes and can it Work on Modern x86 CPU's. Internal Report. Defence R&D Canada. November 2005.

[6] Carbone, Richard. Can Linux be Easily Reconfigured from One Machine to the Next?" Internal Report. Defence R&D Canada. October 2005.

[7] Carbone, Richard. A What to do Avoid Guide in Doing your own In-House Migration. Internal Report. Defence R&D Canada. November 2005.

[8] Michaud, Frederic and Carbone, Richard. Practical verification and safeguard tools for C/C++. Technical memorandum. Defence R&D Canada. Document No. TR 2006-735.

[9] Painchaud, Frederic and Carbone, Richard. Java software verification tools: Evaluation and recommended methodology. Technical memorandum. Defence R&D Canada. Document No. TM 2005-226. March 2006. http://cradpdf.drdc.gc.ca/PDFS/unc57/p527369.pdf.

[10] Weimer, Hendrik. Dissecting Programs. Online article. OS Reviews. September 2006. http://www.osreviews.net/reviews/admin/strace.

[11] Ravi. strace – A very powerful troubleshooting tool for all Linux users. Online article. All about Linux. May 2006. http://linuxhelp.blogspot.com/2006/05/strace-very-powerful-troubleshooting.html.

[12] Cespedes, Juan. Ltrace Linux man page. Man page. Die.net. http://www.die.net/doc/linux/man/man1/ltrace.1.html.

[13] Maurer, Ben. Memory usage with Smaps. Online article. Ben Maurer. March 2006. http://bmaurer.blogspot.com/2006/03/memory-usage-with-smaps.html.

[14] Nguyen, Binh. Linux Filesystem Hierachry. Revision 0.65. Howto guide. The Linux Documentation Project. July 2004. http://tldp.org/LDP/Linux-Filesystem-Hierarchy/html/proc.html.

[15] Burford, Sean. Introduction to Reverse Engineering Software in Linux. Revision 1.26. Howto guide. University of Adelaide. September 2002. http://www.ouah.org/RevEng/t1.htm.

[16] Die.net. ldd man page. Man page. Die.net. http://www.computerhope.com/unix/uldd.htm.

[17] Free Software Foundation. GNU Binary Utilities. Guide. Free Software Foundation. May 2002. http://www.gnu.org/software/binutils/manual/html_mono/binutils.html.

[18] Abell, Victor A. Lsof man page. Revision 4.63. Man page. NetAdminTools.com. http://www.netadmintools.com/html/lsof.man.html.

[19] McDougall, Paul. Microsoft Claims Linux Infringes 42 Patents. Online article. Information Week. May 2007. http://www.informationweek.com/news/showArticle.jhtml?articleID=199501578.

[20] Serverfault. Does an unplugged hard drive used for data archival deteriorate? Online forum. August 2009. http://serverfault.com/questions/51851/does-an-unplugged-hard-drive-used-for-data-archival-deteriorate.

[21] HardForum. Cold Storage Hard Drive Data Life Expectancy? Online forum. July 2010. http://hardforum.com/showthread.php?t=1534173.

[22] BackupAssist. Backup hardware showdown: Tape vs. Disk. Whitepaper. BackupAssist. 2009. http://www.backupassist.com/downloads/whitepapers/marketing/tapevsdisk/BackupAssist_Tape_vs_Disk_Showdown_WP(BA).pdf.

List of symbols/abbreviations/acronyms/initialisms

ACL	Access Control List
API	Application Programming Interface
BSD	Berkeley Software Distribution
C2	Command and Control
CD Co	mpact Disc
CRC	Cyclical Redundancy Check
DDS Digital	Data Storage
DLT	Digital Linear Tape
DMSS	Directorate of Maritime Ship Support
DND	Department of National Defence
DRDC	Defence Research & Development Canada
DRDKIM	Director Research and Development Knowledge and Information Management
DVD	Digital Video Disc or Digital Versatile Disc
Ext2/Ext3 2	nd/3rd Extended Filesystem
FFS	Fast File System
FOSS	Free and Open Source Software
Fsck Filesy	stem Check
GNU	GNU's Not UNIX
GPG	GNU Privacy Guard
GPL	GNU Public License
GUI	Graphical User Interface
HMCCS	Halifax Modernized Command Control System
I/O Input/Outp	ut
IP Intellectual	Property
ISO	International Organization for Standards 9660
IT Inform	ation Technology
LTO	Linear Tape Open
NAT	Network Address Translation

NTFS	New Technology File System
PC Personal	Computer
Petabyte	1 Petabyte = $1x10^{15}$ bytes
PID Process	ID
PKI	Public Key Infrastructure
R&D Rese	arch & Development
RAM	Random Access Memory
RHEL	Red Hat Enterprise Linux
RSH Rem	ote SHell
SCSI	Small Computer System Interface
SSH Secure	Shell
TB	Terabyte; 1 TB = $1x10^{12}$ bytes
UFS	UNIX File System
VLAN	Virtual LAN or Virtual Local Area Network
XFS	eXtent File System

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. OR IGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report or tasking agency, are entered in section 8.)

Defence R&D Canada – Valcartier
2459 Pie-XI Blvd North
Quebec (Quebec)
G3J 1X5 Canada

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED
(NON-CONTROLLED GOODS)
DMC A
REVIEW: GCEC JUNE 2010

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Long-term operating system maintenance: A Linux case study

4. AUTHORS (last name, followed by initials – ranks, titles, etc. not to be used)

Carbone, R.

5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
March 2013	82	22

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical memorandum or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence R&D Canada – Valcartier
2459 Pie-XI Blvd North
Quebec (Quebec)
G3J 1X5 Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
1430JT 15AV34	

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DRDC Valcartier TM 2007-150	

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R) or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

(U) In TM 2006-595, Operating system hardware reconfiguration: A case study for Linux, it was determined through experimentation that a Linux-based C2 operating system can successfully undergo a hardware migration and operating system hardware reconfiguration. The direct benefit of this is the ability to forgo any new operating system reinstallation in order to support newer hardware by using mechanisms internal to the operating system that support changes in hardware. This results in a decreased waiting time for system reaccreditation and redeployment. Since an operating system can evolve over time, it can accommodate changes in the system's hardware, thus presenting a tangible advantage for the Royal Navy, as this allows the operating system to be maintained over the long-term. However, there are complexities involved when maintaining an operating system for long periods. Therefore, this report serves as a ni ntroduction and a simple methodology for performing system maintenance-related tasks that include upgrading, updating, as well as data backups and restoration. This report is neither all-inclusive nor a replacement for qualified system administrators with years of experience. Instead, it can be used as a source of information to provide recommended practices, procedures, and information for helping to plan for long-term system maintenance.

(U) Dans le TM 2006-595, Operating system hardware reconfiguration: A case study for Linux, il a été déterminé expérimentalement qu'un système d'exploitation d'un C2 basé sur Linux peut subir avec succès une migration matérielle ainsi qu'une reconfiguration matérielle du système d'exploitation. Le bénéfice direct est la capacité de ne pas avoir à procéder à une nouvelle réinstallation du système d'exploitation afin de supporter du matériel plus récent et ce, en utilisant les mécanismes internes du système d'exploitation qui soutiennent les changements du matériel. Il en résulte une diminution du temps d'attente pour la ré-accrédi tation et le redéploiement du système. Étant donné qu'un système d'exploitation peut évoluer au fil du temps, il peut donc s'adapter aux changements dans le matériel du système, ce qui présente un avantage tangible pour la Marine royale canadienne. Cela permet au système d'exploitation d'être maintenu à long terme. Cependant, maintenir un système d'exploitation sur une longue période engendre des complexités. Par conséquent, le présent rapport se veut une introduction et une méthodologie simple pour effectuer les tâches reliées à la maintenance d'un système qui incluent la mise à niveau, la mise à jour, ainsi que la sauvegarde de données et leur restauration. Ce rapport n'est ni exhaustif, ni un remplacement pour les administrateurs de systèmes qualifiés avec plusieurs années d'expérience. Il doit plutôt être utilisé comme une source d'informations utile pour fournir des pratiques recommandées, des procédures, ainsi que des informations pour aider à la planification de la maintenance à long terme d'un système.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Free and open source; FOSS; Hardware migration; Hardware reconfiguration; Kernel; Linux; Migration; Operating system; Operating system hardware reconfiguration; Operating system reconfiguration; Patching; Reconfiguration; System administration; System maintenance; Upgrade; Update

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
De science et de technologie pour
la défense et la sécurité nationale

www.drdc-rddc.gc.ca