Achieving the Single Information Environment

Doug Stapleton
Canberra, Australia
dlca2576@bigpond.net.au

Abstract—**The information environment at Australian Defence has traditionally spanned across different networks at different classifications. Bringing these disparate networks together into one Single Information Environment (SIE) is in the way of a grand challenge. This paper outlines the issues involved and plots one possible path towards accomplishing the grand challenge of finally achieving a Single Information Environment. That would mean that Defence could operate on one single network, not multiples as is the case today. It would also mean that Defence users would only have one set of credentials to login to the D-SIE (Defence Single Information Environment, as it will be known for the purposes of this paper).**

Keywords-Defence; information; security;

I. INTRODUCTION

This paper considers four areas that need to be covered in moving to a Single Information environment:

- the challenge of moving from a System High security model to having users at various clearance levels in the single information environment
- protection through encryption of all data at rest and data in transit
- data labeling by classification and new application requirements for re-checking credentials and decrypting data according to user clearance levels
- migration to a new environment, the building sequence of the D-SIE

II. SYSTEM HIGH SECURITY

A. Data Spills

Traditionally, the networks have been separated by classification as the simplest mechanism to prevent data spills. By definition, a data spill is when data is revealed to a person without the necessary clearance. The Information Security Manual (ISM) [3 page 303] defines a data spill as "The accidental or deliberate exposure of classified, sensitive or official information into an uncontrolled or unauthorised environment or to persons without a need-to-know". In the proposed D-SIE, if a person with only a protected clearance was able to see information at the secret or top secret classification, then a data spill would occur.

B. System High Security Mode

As a general rule, all users allowed to access a network have clearance to the highest clearance level of information held on that network. This is known as System High security mode whereby all users have clearance but not necessarily a need-to-know, for all data handled by that system [4] [5].

This effectively ensures that a data spill within the network cannot occur since any user will have the necessary clearance to view that data. Nonetheless they may not be authorised to see the data due to its sensitivity but that is a different issue to a data spill which occurs when a user is given access to data at a higher classification than the clearance level they hold.

C. Multi-level security

This is the nub of the issue in creating the D-SIE. How does one prevent a user from accessing information at a higher classification level than the clearance level held by the user. Effectively every application would become a security enforcement mechanism along with the raw information access given by ordinary file exploration from the desktop (such as by using Windows Explorer). The integrity of the overall system would be compromised if just one system failed to enforce the correct separation of classified information from all users with lower clearances. This is moving into the realm of a multi-level security system. Even if an application correctly enforced the security separation, there are inevitably other paths to the data through direct file access or database queries that could inadvertently allow a user to access data without having the appropriate clearance. This n factor problem rapidly grows to nightmare proportions. Suffice it to say that while data is kept in the clear, unencrypted, that it would be impractical as well as impossible to achieve security certification for the D-SIE on this basis.

III. USERS AND THEIR CLEARANCE LEVEL

A. External checks

The other core issue is that of users and their clearance level. With System High security mode, the work of checking a user's clearance is all done externally before the user is granted the credentials to login to the network in question. Thus all applications on a network can make the reasonable assumption that the user has sufficient clearance to access the application data. The practical result is that applications today are not designed to re-check the clearance of users before providing access to data. Within that overall context, each application generally checks if a user is authorised to perform various functions on the application data. This necessitates users belonging to various groups.

B. Groups

The typical structure is that external security mechanisms are used to meet various grouping requirements by assigning users into named groups. For example, an application may

check if a user is in the AGAO group (Australian Government Access Only), with the user being placed in this group as part of setting up their initial credentials for that network. This places the onus on checking for changes and continued revalidation on the external security processes; which may on occasion lead to inconsistent results on each network, particularly if a person's group membership changes over time without dynamic revalidation.

Some applications check for authorisations by referring to their own security setup of each specific user, giving them access to the application as a general user, super user or application administrator etc.

C. *Revalidating security*

In considering a move to a D-SIE, this implies one full set of all users. The immediate implication is that all applications would need to re-check for the user having sufficient clearance to access the application, even though historically applications are not designed to do this. As a practical example in the D-SIE, consider an application that would have previously had a set of users at the top secret level. Users could be placed into a group for 'TS-Application' that the application could rely on. If a user is in that group they are given access to the application and its data. This would rely on the group membership mechanism being robust enough to prevent any user being added to the group that didn't have the required clearance. This again is problematic and presents a difficulty for accreditation.

A safer approach is for the application to re-check each user to ensure that they have sufficient clearance to access the application and its data. This check against a security database of all users, can control that initial check; DOES USER(X) have (Clearance Level Required); with a yes or no answer. The application would then go on to do its normal authorisation of various levels of user access, but within the context that a data spill would not occur.

IV. Encrypt Everything

A. *Protecting data at rest*

How do we ensure that application access to the data is not given to a user without the appropriate clearance? What is the fail-safe mechanism beyond the grouping and application re-checking suggested above. Even if the application access pathways were all tested and validated, there are other ways to access the data, as a file or via a database query. The only practical way to protect the data at rest from unauthorised access by a legitimate user on the network who lacks sufficient clearance to view the data, is by encryption.

When we think of three separate networks, there are progressively stricter encryption requirements for the various networks as described in the ISM [3 page 236 ff]

However, if the three separate networks are merged into the one D-SIE then all data would be encrypted to the highest security level contained within that data. As a generalisation, the higher classification of the data requires longer key lengths. So the approach changes to using the encryption requirements for TOP SECRET but generating separate keys for each classification.

If all data at rest is encrypted then any user must gain access to the appropriate decryption key by re-checking the clearance level of the user and trying to decrypt the data. Those users without the appropriate clearance will not be able to decrypt the data.

B. *Default classification*

When the application writes data, it must be specified with a default classification. Thus the migration of an application to the D-SIE that previously operated on the TS network would have an application default classification of TOP SECRET. Within the application, if there is reason for data to be written at other than the default classification, then the application must acquire the appropriate keys and encrypt the data accordingly. A practical example might be that a SECRET level source document being manipulated within a TOP SECRET system, would still need to be classified and written to storage with SECRET level encryption.

C. *Overcoming the moat mentality*

In times past, we thought that systems were secure behind the outer walls (firewalls etc.) and that data could be kept in the clear once inside the 'moat'. Cyber attacks in recent years have made it clear that we must now assume that networks are compromised and that intruders can see network traffic and capture data for their own nefarious purposes.

The benefits of an 'encrypt everything' approach to the D-SIE is that a potential intruder can usually only obtain access to data that would be protected by encryption. They must then incur the overhead of decryption before they can make sense of the data; perhaps only to find they have invested considerable time and resources into decrypting a meaningless office memo.

From an intruder's perspective, when all the target data is encrypted; it removes the context, meaning and crucially, the relevance of the data. The intruder or their backers must invest the time and resources into decrypting the pool of data to rebuild that landscape of relevance. That relevance is already available in situations where data is kept in the clear. Who knows whether all that effort will only yield an old copy of football results or other low value data?

D. *Post Quantum Cryptography*

How good is our encryption. Government standards for encryption are now under threat from new developments in Quantum computing which will bring unprecedented changes and allow an intruder to spend reasonable amounts of quantum computing to decrypt data. Current algorithms that would take millions of years of traditional compute power can be broken within minutes using quantum constructs; "By using these algorithms a quantum computer will be able to outperform classical computers by a significant margin. For example, Shor's algorithm allows extremely quick factoring of large numbers, a classical computer can be estimated at taking 10 million billion billion years to factor a 1000 digit number, where as a quantum computer would take around 20 minutes." [8]

In advertising a 2015 *Workshop on Cybersecurity in a Post-Quantum World*, NIST notes that[1]: "The advent of practical quantum computing will break all commonly used public key cryptographic algorithms. In response, NIST is researching cryptographic algorithms for public key-based key agreement and digital signatures that are not susceptible to cryptanalysis by quantum algorithms."

Britain's Government Communications Headquarters (GCHQ), recently called attention to the threat of quantum computing by publishing a paper describing an attempt to build a post-quantum cryptosystem and a quantum attack against this system [7].

"In 1994, Peter Shor of Bell Laboratories showed that quantum computers, a new technology leveraging the physical properties of matter and energy to perform calculations, can efficiently solve each of these problems, thereby rendering all public key cryptosystems based on such assumptions impotent. Thus a sufficiently powerful quantum computer will put many forms of modern communication—from key exchange to encryption to digital authentication—in peril." [6]

The Australian Government has initiated consideration of a post-quantum world with regard to quantum resistant cryptographic algorithms [3 page 241].

So, noting that there will be significant change in cryptographic approaches over the next few years and decades, we can make a general observation on the future attributes of the D-SIE being that:

- All data must be encrypted
- Data is encrypted at the classification level
- There is no clear data in transit across any network
- Data will only be in the clear in memory when being accessed by a computing process on hardware that has assurance provided by a relevant *operating system protection profile*; that data in memory cannot be accessed by any unauthorised user or process.

Ultimately, encryption is a time delay lock in that all encrypted data can be decrypted given sufficient time and compute power. This makes the coming step change of practical quantum computing, whether that takes years or decades, a clear and present danger to our notions of network protection through encryption.

The approach of 'encrypt everything' will bring with it a major change in the design philosophy of applications. Enterprise applications such as 'search' have tended to have their own 'system level' access to data which is then filtered in the final result before passing on to an authorised user. The inherent complexity and performance costs will imply that a new approach of reading a limited subset of data on behalf of the authorised user will predominate; thus avoiding the compute intensive decryption and analysis of data that is ultimately discarded as not being viewable by the end user.

[1] http://www.nist.gov/itl/csd/ct/post-quantum-crypto-workshop-2015.cfm

V. MIGRATING TO THE NEW D-SIE

Can the existing environments be migrated directly to this new vision of the D-SIE. Clearly, the answer to this question is negative, there is no 'change in place' model. There are too many issues to be resolved in changing from the current 'System High' model of security to one where all users share the same environment. From the discussion set out above, it is clear that accreditation issues alone are insurmountable. That means the change must be done by migrating to new greenfields infrastructure where the system can be built carefully according to the new security rules and without the constraints imposed by the prior network and information architectures. Incidentally, that makes a self contained military platform such as next generation warships an interesting candidate, where the system can function with interfaces to the old networks until such time as the whole enterprise can migrate to the new approach.

Migration to the new system requires a migration from multiple separate networks to the one D-SIE. Each application must be well behaved and follow the new rules of encrypting data and always decrypting data on behalf of a user whose credentials are dynamically checked. Applications can keep working data in memory in the clear, but when written to temporary files on that computer/server then they are encrypted with keys specific to that server; not to a clearance level i.e. other servers cannot access that data. This is an important distinction and will guide the implementation of the new system.

We will also need a new perspective on the logging of data in a system where all data is encrypted, including log data. How could information be read from the logs in any case? Security will be about checking for anomalous data and patterns of usage across the encrypted logs. The application systems will need to provide the ability to review their own logs and provide decrypted access to those security personnel with the clearance and need to know. This is a major change in approach that needs discussion.

So we can lay out a sequence for the build of the D-SIE from a greenfields infrastructure point of view.

Firstly the network file storage has a base layer of encryption for the device itself. Above that is a layer of encryption as per the classification of the data. Applications must set a default level of classification so that all data is encrypted, even if it is only to the default level of classification.

Secondly, servers are setup using computing equipment that protects each execution space. A major challenge in a shared user environment is being able to ensure that the operating systems of both servers and end-user devices can keep users separate and ensure that memory allocated to one user or process is not able to be compromised by another user or process on that system.

Assurance of this level of protection is now given by *operating system protection profiles* [1] [2].

To access remote security services, such as determining a user's group membership, an operating system may use a trusted channel, which provides confidentiality and integrity

protection as well as the mutual authentication of the end points of the channel. This may use cryptographic mechanisms or the use of a dedicated physical network to ensure the integrity of the trusted channel.

Security functions may be accessed via remote trusted systems to gain access to centralised security management services for the enterprise.

Thirdly, servers are connected to other devices on the network via encrypted links such that each link uses a different key set to ensure that any devices listening on the network do not see the data in the clear, particularly across insecure networks such as the public internet.

Next the Identity Management system must have real time access to the clearance and nationality attributes of end users, such that a request for access can validate what keys are available to this user for the attempted decryption of data.

The decryption keys returned must be relevant to the individual set of infrastructure (such as different keys on different military platforms). Hence there is detail here on key management issues that are beyond the scope of this paper.

Next applications such as email etc. can be migrated to the new D-SIE, noting that for an application to be accredited on the D-SIE, it must conform to the new security requirements:

- the application must set a default clearance level to which all data is encrypted.
- users accessing data must be revalidated on first use of the application for that session.
- the revalidation allows for the attempted decryption of data using the keys available to the user
- working files for the application and written locally on the server must be encrypted with the server's unique key, making data only accessible to a process running on that same server

This paper references the classification levels used in the Australian Government Security Classification System [9]; but the same principles would also apply to other national systems.

In a multi-level security system, the key issue is how to label data in a meaningful way at the application level. There are applicable standards for metadata, but a consistent application level design approach has yet to emerge.

The AGLS Metadata Standard [10] is an Australian Standard (AS 5044) for cross-domain resource description (see usage of 'protectiveMarking'). In particular, it is intended for information about resources and services on the World Wide Web.

The Australian Government Recordkeeping Metadata Standard Version 2.2 (AGRkMS) [11] describes information about records and the contexts in which they are captured and used. The standard is compliant with the Australian Standards on Records Management (AS ISO 15489) and Metadata for Records (AS ISO 23081). In particular see [11 pages 39, 40 and 76] for security related metadata.

VI. Conclusion

In summary we can see that to achieve the grand challenge of a Single Information Environment requires new thinking and a new approach; all data encrypted both at rest and in transit, a new approach to security logging, applications that follow new security rules, and a migration over time to a new greenfields build of the D-SIE.

References

[1] National Information Assurance Partnership, Protection Profile for General Purpose Operating Systems, verson 4.1, March 2016

[2] Common Criteria (NIAP and BSI), General-Purpose Operating System Protection Profile, version 3.9

[3] 2016 Australian Government Information Security Manual - Controls. http://www.asd.gov.au/publications/Information_Security_Manual_2016_Controls.pdf

[4] DoD Directive 5200.28, 21 March, 1988, page 20

[5] Krutz, Ronald L. and Vines, Russell Dean, The CISSP Prep Guide; Gold Edition, Wiley Publishing, Inc., Indianapolis, Indiana, 2003.

[6] NISTIR 8105, Report on Post-Quantum Cryptography, April 2016

[7] Peter Campbell, Michael Groves and Dan Shepherd, SOLILOQUY: A Cautionary Tale, CESG, Cheltenham, UK, 2014

[8] Simon Bone and Matias Castro, A Brief History of Quantum Computing, Department of Computing, Imperial College, London, April 28, 2000

[9] Australian Government security classification system, version 2.2, April 2015 http://www.protectivesecurity.gov.au/informationsecurity/Pages/AustralianGovernmentSecurityClassificationSystem.aspx

[10] National Archives of Australia, AGLS Metadata Standard Part 1 – Reference Description, Version 2.0, July 2010

[11] National Archives of Australia, Australian Government Recordkeeping Metadata Standard (AGRkMS), June 2015