Computer Security Issues in Open Source Databases

James P Anderson
Box 42
Fort Washington, PA 19034

Introduction

This note is a computer security perspective on the emerging Open Source databases At first look, it appears oxymoronic to talk about security and open source in the same breath However, even though computer security had its origins in attempts to protect 'closed source' (classified) data, there are many points in common, while each has its own unique features and requirements

The paper will not address the unique requirements of the classified world It will attempt to identify the common points between classified and 'open source' data

We note that whether classified or 'open source', the data bases we are considering are generally used in decision support systems of various kinds There are some 'open source' data bases that support accounting systems of one kind or another, from one's personal checking account, to the Federal Reserve system and the entire international monetary system This latter kind of data base, while undoubtedly supporting decisions at all levels, from whether I can afford to buy a new suit, to whether to buy Dollars on the world market to support some trade initiative, is more generally thought of as recording some abstraction of 'value' (wealth?) It measures values of a system (the monetary or financial systems of an individual, company,political entity, or the world), in the same way industrial processes are monitored to be sure that the process under observation is behaving in an expected (and accepted) manner There are undoubtedly finer distinctions to be made than these two, as well as cases where the two kinds of systems overlap sufficiently that it isn't easy to identify the kind of system involved We observe that it doesn't matter because the distinction is made to primarily to illustrate that the security issues transcend the kinds of systems involved

Protection Objectives

The security issues of open source data are predicated on the notion that data is a valuable resource There are significant costs associated with its collection, storage, processing (transforming data into information), and maintenance Therefore, protection of the resource is warranted in economic terms alone Having said that, the question is what do we mean by protection? Put into jargon terms, what is the protection objective?
In the classified world, the protection objectives are reasonably well understood They include protecting the CONFIDENTIALITY of the data In some cases this is expanded to cover how the data is acquired, the so-called 'sources and methods' They also include the objectives of protecting the INTEGRITY of the data as well as the AVAILABILITY of data

While there are some exceptions, (FOUO and 'Unclassified but Sensitive' come to mind), broadly speaking open source data does not have a confidentiality protection objective It does have an integrity protection objective of sorts, as well as an availability protection objective On closer examination of the INTEGRITY protection objective we see that it is part of a larger protection objective of establishing and maintaining VALIDITY of data

In addition to data validity, the question of maintaining proprietary interests is a key problem of open source This particular problem comes in several forms, everything from mutually distrustful parties to keeping hostile interests from the data bases

Data Validity as The Principal Open Source Protection Objective

It is clear that regardless of whether the open source system is for decision support or whether it is for process monitoring (in the broadest sense of the term), valid data is critical to the operation of the system It is argued that without valid data, that a decision support system would devolve to tossing a coin Similarly, with a process monitoring system, if the data reported were not accurate, all kinds of chaos would result from my getting a $25,000,000,000 IMF transfer intended for some country's industrialization program, and the country is credited with my payroll check Note that in a world where money is electronic, the representation of the money (the data) is as good as the money itself Actually depending on one's perspective, it is better that the real thing, since it is like a point is space, occupying no volume, having no dimension, and is weightless

Without getting too fanciful, it is noted that data validity has been a component of data processing for centuries, even when the processing has been manual The accounting system mechanisms of checks and balances, independent audits, multiple signatures on checks, check digits on account numbers and so forth are all measures that were taken to assure the validity of the data representing the business or enterprise being accounted Indeed, the accounting profession is dedicated to determining data validity for businesses and stockholders alike

Data Verification

Even before computers came on the scene, procedures were developed for the purpose of validating data being used by electromechanical accounting systems Older readers will remember when the punched card verification step was an inevitably part of data entry for EAM equipment, and for the first several generations of computers as well The verification step was basically having another person duplicate the data to be sure that both operators produced the input It was assumed that if two operators independently produced the same result, the data was entered correctly Variants on the duplicate the data approach included developing a machine that read the card produced by the first operator and checked each character against that entered by the verification keypuncher If the verifier detected a difference, the machine would stop, and the verifying keypuncher would reproducethe card in error up to the point of the error, enter the correct character, and continue

In the 1960 election, which was the first to predict the result based on early returns from certain precincts, the data was entered 3 times, by three independent operators, and voted on bythe program making the prediction

Multiple Collection Points

This technique has been used in both decision systems and monitoring systems In the latter case, it involves placing multiple sensors in the process (One could have a single sensor at the end of the process that determines whether the process result is good or bad, but that doesn't offer much in the way of data to assert control In WW II, it is reported that the British employed a number of different intercept sites to record German radio traffic Often many the sites would be assigned the same target frequency, and the resulting duplicated transcripts would be forwarded to a central point where the best (in the sense of fewest operator errors) would be used for subsequent processing When reception conditions were bad, messages might be constructed by merging portions of a number of different transcripts The point is that in both kinds of system, decision support or monitoring and control, the replicating observations is a method of providing valid data

Multiple Collection Systems

For some kinds of data, more than one kind of collection system might be involved We might be able to estimate crop production of a given area by using aerial reconnaissance This will provide data (e g bushels/acre)with a certain accuracy We may be able to improve this by interviewing farmers in the selected area asking what their experience has been, given the local climatalogical conditions, whether the weather has been too dry,too wet, and the like Finally, we may obtain other estimates of the yield from other sources, agricultural extension agents, banks, etc Each of these may be considered a collection system The most valid data may be that which is deduced from these many sources by an experienced agronomist

The interesting aspect of these (and many more not listed) data validation steps is that they are steps taken to assure that the best data is available for subsequent use When considered in the light of computer systems, the primary protection objective of the computer with respect to open source data is to preserve whatever validity the data has when it enters the computer

Validity Control in Computer Systems

Physical Security Controls

As in protection provided to classified data, it is necessary to provide a environment that provides protection of the data assets and the logical controls (the operating system and application programs) imposed to preserve the validity of the data The physical controls are imposed provide the foundation protection of the system There exist numerous sources of what constitutes adequate physical controls [1] The physical controls are designed to provide assurance that only individuals authorized by the system owner can obtain physical access to the system

Procedural Controls

Procedural controls are put in place to regulate physical access to data and programs by individuals AUTHORIZED access to the data system The controls recognize the different kinds of access that can occur in the operation of a data system Thus individuals who are OPERATORS have authorized physical access to the data system, and as such can do anything to the data ADMINISTRATORS of various kinds are designated to manage the LOGICAL access to the data systems

The procedural controls used in conjunction with the physical controls provide the physical environment of the data system The environment is that in which the protection objectives are or are not achieved It is evident that without an adequate foundation that any other controls that may be imposed may not have the desired impact due to the weak or non-existent foundation, making it possible for the controls to be bypassed or ignored

Logical Controls

This is where the major effort in computer security has focused for the past 15 years or so The original efforts were guided by experience drawn from dealing with classified information in the paper world It was found that just as the physical and procedural controls provided an environment or foundation for the computer SYSTEM, the computer system needed to provide certain fundamental (foundation) controls for subsequent proper handling of protected data

It turns out that the foundations for handling classified information provide the foundations necessary to achieve other protection objectives including *preservation of data validity*, the protection objective we identified at the beginning of the paper

Before discussing specific methods of preserving data validity in computer systems, a short excursion through what was learned in handling classified data that can be applied in the open source arena is in order

Modes of Access and What They Mean

In discussing security issues in computers, the classified world devised a number of "modes" of operation that covered the main ways in which computers could be used to process classified information depending on the clearances and authorizations of the using population

The modes are

Dedicated - A 'single level' of data, all users cleared and approved for access to that data The security policy (access control) decisions and enforcement done OUTSIDE of the computer system

System High - A range of data security levels from Unclassified to the 'High' of the name of the mode All users are cleared to at least the 'High', Need-to-know applies Again, security policy (access control) decisions are made outside of the computer, except that the computer system is expected to enforce 'Need-to-know' In practice, this means that data bases and files have access control lists (ACL), and individuals are not supposed to be able to access data for which they are not on the ACL The need-to-know decisions are placed in the hands of individuals authorized to use the machine This is modeled after handling of documents and paper forms of classified information Also in practice, this mode is run on machines that are not very penetration resistant As a consequence, the Need-to-know controls may be very weak to virtually non-existent There is no requirement for strong internal controls supported by the operating system

Multilevel - A range of data security levels from a System 'Low ' (It doesn't have to be Unclassified, although it may be) to a High The values of Low and High can vary depending on the specific application NOT all users are cleared for all data processed on the system (Some users may not be cleared for 'High' or anything below, or in the case where the system supports categories of data at a given classification level, not all users may be approved to access all categories of data on the system) For these kinds of systems, the computer enforces a policy that describes who may access what for what purpose (Read, Write, Append, etc) The important thing to note is that the computer is responsible for making and enforcing security decisions based on a security policy encoded into tables identifying what each user of the system is authorized to do The data is distinguished by having labels associated with it that defines an access class to which one or more users have access

The latter mode of operation is the most difficult to achieve safely, and because the computer is being relied on to make and enforce security decisions, the operating system and computer are no long just concerned with application functionality, they are also concerned with security access control (including data movement control) and the protection of the access control decision and enforcement mechanisms in the computer as well as the data bases

The various modes led to characterization of security policies based on who made the decisions regarding who was authorized to see particular data In the System High mode systems, the decisions regarding who is authorized to see data is left in the hands of individual users Each user decides for him or herself who can see data that the user has access to This kind of policy is known as DISCRETIONARY, since it is left to individuals to decide [Note that the data that may be shared is not necessarily owned by the individual making the decision to share it]

The other kind of security policy is one where an organization (business, bureau, Agency, Department, whatever) makes the determination regarding who is authorized access to particular data This is called a MANDATORY security policy because it is expected to be enforced at all times, and not left up to the judgment of individual users This is the kind of policy enforced by a Multilevel system

One may well wonder what all of the operating modes and kinds of policies each can support has to do with maintenance of data validity, the hallmark of open source processing The fact is that without a proper foundation, no method of computer control can be put in place with any assurance that it will be effective The basic hardware andsoftware requirements for processing classified information in the various modes of policy enforcement have been defined in the Trusted Computer Security Evaluation Criteria (TCSEC)[2] There are efforts in Europe to develop similar criteria [3], and in the United States [4] to develop an updated set of criteria that emphasize requirements for processing unclassified data of the sort discussed in this paper

Hardware Integrity

There are two aspects to the hardware integrity issue Those measures built into a system to detect (and in some instances, correct) errors in data that would destroy the validity of data, and those measures built into hardware systems that protect the operating system, and application programs from logical tampering

The first set of measures includes use of parity checks on internal memory (RAM) and storage devices (Disks) Recent technology has given us RAID systems that provide for writing data twice in two different locations on large cheap disks as a means of maintaining data validity, when it is written to storage Error detecting codes (CRCs) computed for data stored, and error correcting codes are also elements of hardware integrity that contribute to the maintenance of data validity Most modern systems have some form of checking capability to apply to data in storage and data in movement across networks

The other aspect of hardware integrity is the structure of the operating system and the hardware features available to support operating system features affecting data validity These include the procedura' features of continuous protection against tampering or unauthorized changes to hardware and software, provision of a separate domain for the operating system security components, mechanisms (hardware and software) to periodically validate the correct operation of the system hardware and firmware (if present), ability to isolate resources so they are subject to access control, and separate address spaces for each user,

Additional steps to support integrity for data validity are usually implemented in the applications software itself This paper will not attempt to enumerate all of those approaches proposed, tried or routinely used For further details on integrity measures see the papers by Mayfield [5,6]

Proprietary Interests

Much of the open source data is collected, and packaged as part of an on-going profit-oriented enterprise The data is sold as a commodity Unlike most commodities data is never exhausted after it is delivered to a consumer, it is self-replenishing It is not destroyed or physically remvoed from one place to another Further, once data leaves the hands of its 'owner', it is no longer possible to exert any physical control on that data As a consequence, the data is susceptible to unauthorized copying regardless of the form it is in

Technology can be applied to assure a data owner that data he has sold can only be access by a given buyer However, there is no practical technology that prevents a buyer from making unauthorized copies for his own convenience and use or even for 3rd parties Onemerely has to observe the practices ofsoftware pirates, video pirates, book pirates etc to see the phenomenon No matter how elaborate a technical scheme, one constructs to prevent unauthorized 3d party distribution, nothing is foolproof This is because it is necessary to render the data intelligible at some point in order to make it useful

The best one can hope to do is uniquely identify each copy of a data commodity in order to identify the SOURCE of a pirated copy if one shows up Here the ability to digitally 'sign' documents using public key or crypto sealing techniques makes it possible to bind a 'fingerprint' (e g document serial number) to a document in such a way as to uniquely identify any document [7] The only question to using the technology is how small an entity in what kind of setting (print, electronics) can one afford to tag OF course, it doesn't PREVENT willful disregard for property rights, it merely makes it possible to identify a possible culprit

Confidential Interest

The maintenance of confidential interest has several dimensions everything from hiding direct interest (identity) to hiding a specific interest (confidentiality of intent) To some extent the hiding of direct interest is one that has been around in less open arenas for some time Use of cut-outs of various kinds and sophistication has been used for example to disguise a corporate interest in acquiring another company To the extent that direct interest needs to be hidden, go-betweens are well understood

It is not far fetched to want to disguise interest in some topic One does not have to be the CIA or even a government entity to want to hide interest in a particular set of data in some data base A common example would be a business looking to acquire another company The desire to conceal interest is driven not wanting to perturb the price of the stock by the interest shown in the business, thus making a potential acquisition that much more expensive

The problem of hiding specific interest is somewhat more complex The provider can monitor all requests and if the problem is important enough, could play the requests against his product to see what people were interested in [One might even attempt to justify such activity as 'Service Quality Monitoring' or 'Market Research']

Hiding specific interest can be accomplished by designing questions to include the specific items among a broader set of items If data is available from two or more sources, the questions might be arranged such that some of the sought after data is received from one source, the balance from a second (third) The techniques of how to include items of specific interest in broader questions is situation dependent However it is worth while noting that it can be done For more information on the nature of the problem, see Chapter 6 in Denning [8]

Conclusions

Just because one is dealing with open source, doesn't mean one has no computer security interests The question of data validity alone puts
as much emphasis on system design and implementation techniques as does data confidentiality in another setting What is unknown for the integrity question is how much assurance is required for data validity mechanisms For any of the decision support systems, it is suggested that the features and many of assurance requirements of the C1, and C2 class of trusted systems are a necessary foundation for the applications level {functions/mechanisms} used to maintain data validity While some of the proposals for implementing integrity controls appear to be radical departures from "traditional" computer security architectures, on closer examination, they are mostly application of specific techniques that require the TCSEC C1 or C2 foundations for self protection[1] if nothing else

[1] The needs of the original computer security work were advanced with the notion of a Reference Monitor as a computational abstraction for secure (confidentiality preserving) computing The reference monitor idea was accompanied by 3 properties
1 It must be tamper proof
2 It must always be invoked
3 It must be small enough to be subjected to analysis and tests to insure its completeness (and correctness)
The C1/C2 foundation requirements address these points, particularly the self-protection aspect

References

1 NIST FIPS PUB 31 *Guidelines for ADP Physical Security and Risk Management,* November 1974

2 DoD 5200 28-STD, *"Department of Defense Trusted ComputerSystem Evaluation Criteria,"* December 26, 1985

3 *"Information Technical Security Evaluation Criteria, Provisional Harmonized Criteria"*, Version 1 2, June 1991

4 USITSS

5 Mayfield, et al C Technical Report79-91, *"Integrity in Automated Information Systems,"* NSA, September 1991

6 Mayfield, et al IDA Document D-967, *"Integrity -Oriented Control Objectives Proposed Revisions to the TCSEC"*, October 1991

7 Merkle, Ralph C , *"Protocols for Public Key Cryptosystems"*, Proceedings of the Conference of the IEEE Technical Committee on Security and Privacy, May,1980

8 Denning, Dorothy E R , *Cryptography and Data Security,* Reading MA, Addison-Wesley, 1982

FIRST INTERNATIONAL SYMPOSIUM: NATIONAL SECURITY & NATIONAL COMPETITIVENESS: OPEN SOURCE SOLUTIONS Proceedings, 1992 Volume II - Link Page

Previous	OSS '92 John Gage, SUN Microsystems, Open Sources, Open Systems,
Next	OSS '92 Erika D. Nobel, From A to Z: What We've Done with Open Sources,

Return to Electronic Index Page