

A Multi-Levelled Approach to Intrusion Detection and the Insider Threat

Rita M. Barrios

Computer Information Systems, Cyber Security, University of Detroit Mercy Detroit, Detroit, USA
Email: barriorm@udmercy.edu

Received November 8, 2012; revised December 15, 2012; accepted December 26, 2012

ABSTRACT

When considering Intrusion Detection and the Insider Threat, most researchers tend to focus on the network architecture rather than the database which is the primary target of data theft. It is understood that the network level is adequate for many intrusions where entry into the system is being sought however it is grossly inadequate when considering the database and the authorized insider. Recent writings suggest that there have been many attempts to address the insider threat phenomena in regards to database technologies by the utilization of detection methodologies, policy management systems and behavior analysis methods however, there appears to be a lacking in the development of adequate solutions that will achieve the level of detection that is required. While it is true that Authorization is the cornerstone to the security of the database implementation, authorization alone is not enough to prevent the authorized entity from initiating malicious activities in regards to the data stored within the database. Behavior of the authorized entity must also be considered along with current data access control policies. Each of the previously mentioned approaches to intrusion detection at the database level has been considered individually, however, there has been limited research in producing a multileveled approach to achieve a robust solution. The research presented outlines the development of a detection framework by introducing a process that is to be implemented in conjunction with information requests. By utilizing this approach, an effective and robust methodology has been achieved that can be used to determine the probability of an intrusion by the authorized entity, which ultimately address the insider threat phenomena at its most basic level.

Keywords: Bayesian Belief Network; Database; Insider Threat; Intrusion Detection

1. Introduction

As far back as the 1970's, detection of a data breach at the database level by an authorized insider (also known as the insider threat), has been an issue that has plagued the information technology community. The research put forth addresses the insider threat issue by presenting a multileveled approach to Database Intrusion Detection.

1.1. The Insider Threat

The theft and exposure of the critical data components that resides in a relational database by the authorized insider is on the rise [1].

An authorized insider can be defined as an individual who has been granted privileges to utilize or modify the critical data components. This entity can be characterized as an entity that chooses to abuse their role to perform malicious activities. It is because of this type of threat that has given visibility to the need for an automated solution that enables detection of this type of breach [10,20]. While the insider threat presents a trust issue that cannot be solved with this research, the framework as

presented did, however, aid in the reduction of exposure when the motivations and subsequent actions of the trusted user can no longer be relied upon.

Finding those trusted entities that are capturing the confidential data components is a task that is difficult at best [1,2]. Identification can encompass a complex decision making process on several levels of abstraction which includes having an understanding of the daily non-threatening, functional actions of the users as identified in the usage logs as well as the measuring those actions against the defined access control policies associated with the related user classification [3]. The challenge within this scope becomes how identification can be successfully accomplished to distinguish between the dynamic and valid usages of the data components vs. the abuse situation. With the introduction of an automated and robust extension to the current research in database Intrusion Detection Systems (IDS), it is possible to overcome this challenge.

With this challenge in mind, the research presented focused on a novel combination of established methodologies in data mining, policy identification and abuse

identification in an attempt identify inappropriate behavior as demonstrated by the authorized users of a relational database. Additionally, there are two secondary objectives presented in this research. The first being the creation of a supervised learning component of the presented Database Intrusion Detection Systems (dIDS) to determine the validity of the “normal” behavior and the latter is to develop the definition of the behaviors that were needed to identify, classify and respond to the introduction of new transactions. These objectives are realized by the pairing of the defined access control and security policies with the usage behaviors as found in the database logs.

1.2. Relevance

When the theft of critical data components is successfully executed by the authorized insider, corporate trust begins to deteriorate among its consumer base as well as exposes the organization to various legal issues due to the violation of local, state and/or federal laws and regulations.

This phenomena occurred recently when [4] reports that the data breach of 2008 at T. J. Maxx is expected to realize costs of more than 10 times the original estimate to a record \$4.5 billion. This is about \$100 per consumer record for each of the 45 million credit card account numbers stolen over an 18-month period [4]. It is expected that the cost associated with the breach to lower profits by \$118 million in the first quarter of FY 08 [4].

The insider threat is considered significant, since there have been many cases presented in the literature where a security breach had been successfully accomplished by the actions of an internal user when there has been an advanced level of authorization granted. Statistics as presented in a study by [1], which had been conducted by the FBI in 2006 identified approximately 52% of the respondents had reported unauthorized use of information resources by internal users along with 10% of the respondents unsure if they had been exposed. Further, judicial proceedings as presented by [5] as well as is documented in industry publications as presented by [4] suggests that there is a significant degree of loss to those victim organizations where there has been an exposure of data as a result of inside. Reference [5] also presents a case where a senior database administrator (DBA) had pled guilty to stealing 8.5 million consumer information records over a five-year period, which subsequently had been sold for approximately \$580 k. Reference [5] further exposes the problem of the insider threat by presenting a case where the directors of admissions and computer center operations at a Manhattan college were indicted on charges of fraud after setting up an operation where people who had never attended the college paid

between \$3 k and \$25 k to obtain forged academic transcripts. As can be seen, the anomaly of the insider threat is consider a cross-cutting concern as it is not restricted to a single industry but can encompass all types information systems.

Many researchers tend to focus on a single aspect of the overall solution to database intrusion detection. It is not clear as to why the merging of the two critical components of transaction validation and abusive activity validation has not been attempted to create a well-rounded and complete solution even though many researchers do recognize the need for both components [2,3,6-8]. With this obvious omission, it appears as if this type of multileveled solution may in and of itself, pose a task that is too difficult to accomplish.

1.3. Guide to the Paper

The research presented takes on the following progression. Section 2 presents a brief but thorough review of the literature associated with research on database intrusion detection, digital rights management and data mining for behavior patterns. Section 3 presents a brief account of the methodologies that were employed to successfully implement the framework while Section 4 presents the post implementation findings. Section 5 concludes this work with a summary of research presented along with an examination of future works to be achieved.

2. Prior Research

2.1. Access Controls

Although advancement in database access controls has made significant progress towards securing data that resides within the database, there are still limitations of how much can be prevented when considering the insider threat phenomena. This becomes apparent given that the majority of the strides made thus far are focused on addressing the functions of proper authorizations. Since the insider is already authorized these methods will not prevent the theft and/or exposure of the data components [9].

2.2. Intrusion Detection Methodologies

Intrusion Detection Systems (IDS) have been a focused research subject for decades with significant attention given at the start of the 21st century. It was at this time that [10,11] as well as others began to present the foundational concepts in formalized research and publications. Because of the clarity presented in the [11] research, the discussions presented in the following paragraph are centralized around this work however; the seminal works are identified as well as the basis for [11] study. It should be noted that the research presented in the [10] study

follows the same presentation as [11] whereby the seminal works form the foundational concepts.

IDS can be defined as a system with a goal to defend a system by raising an alarm when the protocol detects that there has been a security violation [11,12]. With this in mind, [10,11] as well as the seminal works in intrusion detection, identify two primary principles in the IDS model, anomaly detection and signature detection [12,13]. Anomaly detection is defined as flagging all abnormal behavior as an intrusion [10,12-14]. While Signature detection is defined as flagging behavior that is relatively close in comparison of some defined, known pattern of an intrusion that has been previously defined to the IDS [10-13].

2.3. IDS Implementation

When implementing IDS the focus must be on the three primary components: the audit data, detector model and output that will be used in the follow-up process. The detector model and its underlying principles is the primary component of the system [10-19]. Additionally, IDS protocol is further categorized base on the type of intrusion recognized by the system. These categorizations are defined as follows: The Well-Known intrusion that identifies a static well-defined pattern discernible with the intrusion being executed in a predictable fashion. The Generalized intrusion is similar to the Well-Known category but is variable by nature. Lastly, the Unknown intrusion is identified by a weak coupling between the intrusion and a system flaw. This intrusion category is the most difficult to understand [5]. It is this category, the Unknown Intrusion that is the focal point of the presented research when coupled with the insider threat phenomena.

2.4. Database Intrusion Detection Systems

Historic as well as current research in the area of IDS and access control methodologies does not support the identification of intrusions at the database level [3,6,7,16]. As can be seen, most notably in a study presented by [11] on threat monitoring, there are three classes of database users: 1) The masquerader who has gained access to the system by impersonating an authorized user; 2) The legitimate user who misuses his/her authority; 3) The clandestine user who is operating in stealth mode and nearly undetectable. As [11] notes, the legitimate user can be most difficult to detect using standard audit trail data, as abnormal behavior is difficult to detect when the occurrence is minimal or when the standard rules of operation are often subject to exceptions and modifications. If the access control methods are the focal point, a misconception that simply having the right levels of access control applied to the data components as defined by [16] is suf-

ficient to protect the data or that these methods will function as a deterrent to abusive access behavior. However, given statistics as noted previously on data breaches involving the legitimate user, it is understood in the industry that standard IDS with standard access control is not enough to prevent the insider threat risk. When taken in context with a legitimate user, these forms of detection and access control often fall short in detecting the abusive behavior [3,6,7,16].

Often, the developers of the common IDS make a false assumption that an entity accessing critical data is authenticated and authorized via external supporting automated security measures. These external authentication/authorization methods reside primarily at the networking and/or operating system level. Another assumption is that authorization and authentication is a failsafe and always successfully identifies the entity as behaving in a trustworthy manner based on the high user database level access [6,7].

To aid in the closure of the gap between access rights, information protection and levels of responsibility, research in the area of the authorized insider threat risk as it relates to the capture of the critical data components has started to take shape. As seen in the [1] study an attempt is made to identify the person making a request for information via usage of the DBMS audit logs to determine whether the requestor is functioning within the boundaries of their security capacity. Again, as presented by [8] an attempt is made at identifying suspect actions via the usage of a quantitative measurement of transaction violations as had been mapped from the Database Management System (DBMS) audit logs. This measurement determines whether the requestor is making a "legal" request [8]. However, both of these studies fall just short of the proactive, dynamic and automatic identification of an abusive use of the data components by the authorized insider.

As previously noted, [20] attempts to identify the unauthorized insider by constructing a Networked Bayesian Network (NBN). This was done in an attempt to project the probability of an intrusion when critical data components are linked within a transaction. The deficient factors identified in this study are the authors' base assumption that 50% of all insiders are attempting to breach the system, leading the reader to conclude that the assumption is an unrealistic expectation given that there are no documented references in the study that lends itself to this percentage. In addition, the authors [4] realize that their proposed method is ineffective when applied to the authorized insider threat risk. However, if applied properly and with additional controls such as the use of the corporate access controls and security policies, this deficiency may be resolved. In addition to the inability of being able to identify an intrusion when executed by the

authorized insider, the current research as demonstrated above produces a mutually exclusive view on abuse identification of the authorized user and the identification of a potentially harmful transaction when in actuality the two components should be working to compliment one another's strengths [1,6,7,8,20]. Therefore, the research put forth begins to fill the gap that exists between the concepts of malicious transaction and abusive action identification by expanding these concepts to incorporate the defined access control, security policies as well as the behavior of the authorized user to identify a viable, proactive solution that is both dynamic and automated.

3. Methodology

The foundations of this study were focused on three primary facets. The first was the research proposed by [21] with their methodologies of mining association rules within a large set of data using the Apriori Hybrid Algorithm. The methodologies put forth by [14] in the area of utilizing the Stochastic Gradient Boosting and the Bayesian Belief Network algorithms to determine probabilities was the second pillar for this study. Thirdly, current methodologies utilized in the dynamic maintenance area of security policy, known as Digital Rights Management (DRM) served as the final pillar to secure the foundation of this research. Along with the novel approach to database intrusion detection that guided the presented research, a series of modified Intrusion Detection System heuristics has been presented to provide a solid foundation for the acceptance of the results of the approach. The following paragraphs outline how these objectives were achieved.

3.1. Development Approach

The process began with the Trusted User initiating a transaction. Within the context of this study, whether the transaction is initiated via internal or external means has not been addressed but is considered in future work. Once the transaction has entered the presented dIDS system, several processes were initiated to determine the probability of an intrusion. The results from the probability assessment were stored in the dIDS repository for future reference by the dIDS. At this time, the transaction continued on to its completion since this study was focused on intrusion detection and not intrusion prevention. Future work focuses on extending the work presented into the intrusion prevention research area within the context of the database environment.

3.2. Association Rules

Following the building of the data repository that housed the various dIDS signatures, generation of the dIDS

training data dependencies signatures was the next logical step. As in most organizations, certain data components are dependent upon other data components. Often usage outside of the normal transactional scope is an anomaly in and of itself. For instance, the retrieval of only a consumer name data component really does not garner one much information however if that same consumer name component is used in tandem with the corresponding consumer address component, one can make inferences about the data that was retrieved. It then makes sense if the selection of certain data components is without their complimentary counterpart, one can reasonably conclude that an intrusion may be occurring. However, obtaining the data dependencies can be a daunting task given that there may be thousands of various data component combinations open to selection. As [22] have reasoned, often times, in intrusion detection, the training data is developed utilizing a significant amount of expert information about the system and often times this domain knowledge is difficult to obtain. Continuing with the implementation of the [21] Apriori Hybrid Algorithm to mine the association rules, the difficulty of the domain expertise is greatly reduced. Following the acquisition of the training data, it is then subsequently used to identify known intrusions. As [5] point out there is a direct relationship between the quality of the intrusion detection model and the quality of the training data obtained thru various data mining techniques. With this quality concern in mind, the study presented in this research utilized an Apriori Hybrid Algorithm in order to determine the most appropriate data dependency signatures (or rule associations) since there can be a high correlation between the combinations of the data components. Utilization of this type of algorithm is common during a data mining operations when obtaining a selection of relevant facts (data components) where the members have a high degree of correlation [21]. Application of the algorithm in this regard helped to limit the creation of data component combinations to only those where the historical pattern has been consistently demonstrated throughout the data. Once developed, the application of the algorithm to the TPC-C data occurs in order to determine which data components appear to be of significance for the processing environment of the prototype dIDS. Since the creation of the training data is a generic process, utilization of the algorithm to process the historical data in order to determine the data component significance is appropriate. As such, it is also appropriate to apply the training algorithm to a variety of input data following the definition of the components within the pool of information to develop newly identified association rules.

To begin this novel approach to database intrusion detection, an unsupervised learning process was initially

employed in a data-mining environment to establish the baseline rules, which developed the data association rules that established the behavior correlations. Rule associations algorithms are well researched as noted above. These methods are considered the standard in data mining when trying to establish data correlations. The Apriori Algorithm as implemented by [23] is said to be the most popular of these types of mining operations [24]. However, an extension to this algorithm as developed by [23], called the Apriori Hybrid was utilized in the development of the initial data signatures because of its wide acceptance within the data mining community. This is due in part to the algorithm's quality levels when mining user behavior, patterns of access and the assigned classifications from historical data [24]. To implement the Association Rule algorithm, two steps are taken into account to satisfy the user-specified minimum support and the user-specified minimum confidence in parallel [23]. These two steps are as follows:

- Apply the minimum support to find all frequent item sets in the database
- Form the rules using the frequent item sets as defined in the first step and minimum confidence constraint.

Typically, the first step is more challenging since it involves searching all item combinations. Given the growth rate of the item set can be expressed as a potential for exponential growth depending on the number of items in the item set, a method of deterring this growth can be found with the implementation of the downward-closure property of the support constraint [23,24]. This property guarantees that for a frequent item set, all of its subsets are frequent and therefore for an infrequent item set, all of its supersets must be infrequent.

Since the Apriori Hybrid Algorithm exploits the best features of both the Apriori and the Apriori TID in addition to being one of the most popular of the Association Rules algorithms as noted by [24], its foundational properties have been employed in the presented research.

3.3. Probability of Intrusion

As evidenced in historical and current research, fuzzy logic and/or neural networks have been successfully used to determine whether an intrusion has been encountered [3]. Since "normal" behaviors are often known within the data processing environments and patterns of behavior can be established from the historical information, the utilization of a similar approach to the neural network IDS solution can be implemented utilizing the more defined decision tree methodology to determine the probability of an intrusion.

The data gathered during the data mining process as outlined above was then utilized to refine the prototype system by utilizing the supervised Stochastic Gradient

Boosting decision tree process to establish the probability of whether a given signature as created by the Apriori Hybrid Algorithm is considered an intrusion [14,25].

It should be noted that the recommend practice as suggested by [25], is to perform both the Stochastic Gradient Boosting tree creation as well as a single tree. This is because the Stochastic Gradient Boosting method is more like a "black box" methodology that is highly accurate however; it is difficult to visualize the relationships established during the process [25]. This recommendation of running both methods has been followed in the presented research to ensure the most complete information can be realized in building the model's accuracy as well as being able to fully understand the relationships.

Once the prototype had been successfully built with the association rules based on the Apriori Hybrid Algorithm as well as the detection signatures as identified in the Stochastic Gradient Boosting methodology, this same learning process was employed to account for new entities making requests for information. Upon the discovery of a new entity, the behavior signature repository was updated accordingly with relevant data.

3.4. Current Security Policies

In most organizations, the ability to dynamically create and maintain acceptable use policies tends to be an extensive and resource intensive process. For most, the development life cycle goes thru an iteration of steps ultimately ending up at the point where the policy must be published [12]. Today, most organizations' view on publishing the policy's modifications requires the updating of web pages, hard copy documents as well as applying any needed updates to the information systems via physical code modifications. Since a process so resource intensive can take weeks or even months to realize, often times, it is near impossible to determine if a violation of policy has occurred until sometime later. In recent times, the application of digital rights management systems (DRM) to allow for policy development and distribution is taking shape [26]. A DRM system is a system that allows for the management of the actions and entity can perform on a digital resource (the data) as well as controlling the usages of the resource within the information system. As [26] notes the ability to specify and manage the rights of an entity is one of the most important features of the DRM. Unlike standard authorization mechanisms, the DRM is meant to give specific rights to specific entities for a specific amount of time [26]. Bringing this notion of the DRM into the context of this research by building a DRM-like repository allowed for dynamic, real-time policy development that can be accessed at will by the presented intrusion detection system.

3.5. Detection of Abusive Activities with the Bayesian Approach

Since the Bayesian Belief Network (BBN) methodology employs a reasoning mechanism that enables the determination of the probability of an event occurring when various factors are present, usage of the BBN is the most effective method of detection for this study.

The usage of BBN has been supported in recent literature as a viable method of intrusion detection. In [27] study, the Bayesian approach combined with a visualization component is defined to create an interactive intrusion detection system in an attempt to reduce the number of false positives presented in current intrusion detection systems. Again, in the [28] study, the Bayesian approach is applied to improve the effectiveness of the detection mechanism in the presented intrusion detection system for a mobile ad hoc network. Reference [20] used the Bayesian approach to expand the independent environment variables often present in intrusion detection to propose a networked Bayesian Network to understand the correlation between these environmental variables which may be used to identify an intrusion within a relational database. Therefore, with the objectives of this

research endeavor in mind and in keeping with current research, the probabilistic approach of utilizing the BBN where the conditional probability and the causality relationships between the variables as defined have been applied to the presented intrusion detection system. Since the diDS does have knowledge of the acceptable behaviors, relevant security policies as well as the data dependencies, the BBN can make a reasonable assumption of the probability that an intrusion has occurred even when presented with new information.

Implementation of the BBN took the following path in the research presented. Initially, the data mined as described in the preceding paragraphs was used to identify the variables (nodes) of concern for the detection model as well as their association rules. With this information, the Directed Acyclic Graphs (DAG) was developed to visualize the conditional properties of the relationships presented. Following the creation of the DAGs the probabilities for each node was developed.

3.6. Information Flow Overview

The flow of information within the presented research and identified in **Figure 1** is as follows which is similar




Figure 1. Information flow diagram.

to the [20] study. Each of the contributing factors, which include environment, policy as well as data component combinations were given a conditional probability. If the computed probability of the information request, (given the identified variables), fell within an acceptable range, the transaction was identified as not being an intrusion. However, if the probability fell outside of the acceptable range, the transaction was defined, as a potential intrusion where by further steps were completed in an attempt to make a determination as to whether or not the “true” intrusion flag is accurate. These further steps include looking into the system to determine if there are any new applications as well as users accessing the system. Should this situation present itself where the dIDS encounters a “new” entity, the signatures of the entity were stored for utilization in future detection processes. Should the system identify an actual intrusion, in addition to storing the intrusion signatures, the transaction was flagged as an intrusion that should be further examined.

It should be noted that in order to determine the most optimal thresholds, various probability thresholds were utilized. This aided in the determination of the level that the presented research is most effective with the goal being that the model should employ the most restrictive threshold possible while maintaining an acceptable level of detection.

3.7. Management of New Information

The research as is presented does not suggest that the flagged transactions should be prevented from completing its processing cycle. This decision should be left the system administrators given the business model in which the database operates. Future work focuses on extending this research into an intrusion prevention model to facilitate the prevention of potentially harmful transactions from reaching the database.

Continuing on to the acceptance of the new information within the dIDS, “new” data was collected via the development of a simulated payment processing application, where the necessary data components were established and persisted in a relational database. In addition, this subset of “new” data was utilized to develop the intelligent features needed to enable the system to automatically identify the new entity and “learn” from that new relationship for future reference.

To begin, a SQL Server 2008 Enterprise Edition database was created modeling the TPC-97 data structures. Give that one of the objectives of this research was to implement the Database Intrusion Detection System (dIDS) in an environment that supports the very large database (VLDB) concepts, the TPC-97 was selected as the data model of choice given its ability to support large

amounts of data and transactions.

In order to generate a significant amount of data for the TPC-97 model, a data-generator software product was purchased under a research license via Red Gate Software, Ltd., a private limited company located in the UK. This software, called Data Generator, was specifically developed for the SQL Server environment to generate meaningful test data. To support this functionality, the generation of 1 million customer accounts as well as supporting order entry information, which included the customer order history using the primary keys of the schema’s table object occurred. Following the build of the baseline TPC-97 objects, the specific tables for the intrusion detection system were generated along with the incoming transactions. The policy and entity information was manually entered using standard INSERT statements of the SQL query language given that the data was not complex in nature. To support incoming transactions to the TPC-97 database, 925 random transactions were generated based on the policies of the database. Following the requirements definition of the experiment the Apriori Hybrid algorithm was implemented in the Java programming language and applied against the TPC-97 database to determine the most common patterns of data contained within the data. These patterns served as the baseline patterns used in the dIDS.

To begin the process, the generated transactions were sequentially read into the dIDS where an attempt was made to seek out any existing patterns previously identified as well as any policies that may be set in place to govern the transaction processing. If the policy and pattern were found within the signature repository, the transaction was identified as being proper. However, if only a portion or no information was gathered about the transaction, it was subjected to further scrutiny by first going thru the Stochastic Gradient Boosting algorithm to understand the probability of an intrusion based on the transactional statically defined components such as the data relationships and defined policies. If the transaction continued to be identified as an intrusion by exceeding the pre-defined threshold, the transaction was then subjected to the Bayesian Belief Net algorithm to determine the probability that an intrusion was occurring based on the uncertainty of the environment and the prior behavior of the requestor. It should be understood that the requestor could be either a human or an electronic source. Should the transaction continue to be identified as an intrusion, it was flagged as an intrusion then logged as such in the signature table to be used in the analysis of future transactions. Manual verification and validation was then applied to determine whether the transaction was a true intrusion or classified as a false positive.

To support the learning component, if the transaction was not classified as an intrusion, it was also logged in

the signature table along with the information to identify the transaction as a potential new signature. Like the intrusions that were identified during the probability of intrusion calculations, a manual verification and validation process was applied to determine the validity of the new log entry.

Since the availability of a standardized set of metrics used to measure intrusion detection at the level of the database are not readily available, the NIST algorithms that are used to measure network intrusion detection systems were modified to consider the database environment.

Instead of using the network packet as the focus of the measurements as is done in network intrusion detection systems, this research focused the attention on the incoming transaction that are requesting database services. This insider entity was defined to take on the role of an authorized application user, a Database Administrator or an automated program that requests database resources. These roles were defined in the additional database objects that were built to support the intrusion detection system. No programmatic distinction was made during the development of the functional dIDS as to which entity was requesting the information meaning that the policy definitions were the driving force in the identification of the requesting entity.

4. Findings

Intrusion detection at the database level remains a new concept within industry and research. Attempting to find a relevant study to measure the presented research against was difficult. That said, the [5] study was selected as a measurement due to similarities in methodology. In the [5] study, the authors built user profiles of normal behavior as a baseline to detect anomalies. In addition, included in the [5] study was the deployment of the Bayesian approach to estimate probabilities to strengthen their intrusion detection process. Given these commonalities, the information presented in this study served to create the goal measurements for the presented work. Other studies on database intrusion detection identified in this research had a significant enough difference in approach and methodology that they could not be used to create a reasonable and like objective value.

The metrics identified in [7] are presented at a degree of abstraction higher than what was useful for this study. Therefore, the data identified in [7] was expanded and mined in an effort to make visible enough natural detail to provide a reasonable comparison to the presented work. Given that in [7] the successful detection rate was identified at 38.38%, it also served as the baseline for successful detection in this study. The presented study achieved similar results that fell within the range of

<61.62% for false responses and $\geq 38.38\%$ positive responses.

Additionally and where applicable, specific baseline measurement values identified in [7] were adopted and included in accordance with the Information Assurance Directorate (IAD)-US Government Protection Profile for Intrusion Detection Systems as described in [7,28,29,31]. Specifically, these measurements inclusive as qualitative measurements that relate directly to the accuracy of the detection model as defined in the following.

4.1. Coverage

Coverage (c) is determined by the rate which IDS can successfully identify an attack under ideal conditions. The original measurement as identified by [29] is concerned with detecting and measuring the number of unique attacks. This measurement, in a signature-based system, is achieved by verifying the number of valid signatures (s) and mapping them to a standard naming schema. The entering transaction (t) is then measured against the known signatures to determine if the transaction exists within the signature baseline. If the transaction does exist, it is then considered a valid request for information since it meets the criteria of an authorized entity. If the signature does not exist, it is considered a successful recognition of an anomaly and counted in the measurement as noted in Equation (1). To support database intrusion detection in regards to the authorized insider, the measurement presented was adapted to focus attention on the rate by which IDS can identify an anomaly that has been initiated by authorized insiders at various levels of authorization. Therefore, only transactions that are initiated by a trusted user are considered in the measurement. This measurement achieved a rate of $\geq 38.38\%$ as is indicated using the data presented in [7].

$$c = \sum(t \notin s)/s \quad (1)$$

4.2. Probability of Detection

This measurement determines the detection rate (d) of attacks correctly detected (b) by an IDS in a given environment during a particular period (T) in minutes [29]. As with the Coverage measurement, probability of detection assumes that various types of attacks were measured. Given that this study is focused on one type of attack, the measurement was adjusted to focus on the various levels of authorization of the information requestor as opposed to the various types of attacks. Given that the false positive rate is directly related to the detection rate, care must be given to ensure that the scenarios are exact and consistent when used in both measurements. This measurement, in accordance with the baseline data set as presented in [7] achieved a rate of measurement of $\geq 38.38\%$

of correctly identified attacks.

$$d = \left(\frac{b}{t} \right) * T \quad (2)$$

4.3. Volume of Data

This measurement determines the difference in the volume of data (v) the dIDS can manage when presented with a large mass of transactions as compared with the pre IDS implementation. While data as presented in the NIST IR-7007 identifies the number packet/second for the network IDS, when put into the context of this study, the transaction will be the unit of measure. Additionally, this measurement will identify the change of data volume pre and post dIDS implementation in order to identify any latency issues that may be present.

Given that this measurement is very subjective to the system environment, the initial transaction sets were processed without the implementation of the dIDS to capture the rate of processing using the maximum (m) volume of data. This measurement was considered the baseline processing rate. Following the establishment of the baseline measurement, the same transaction set was processed through the implementation of the dIDS (i) to determine the change in processed transactions through the dIDS implementation. The pre and post IDS expressions used to calculate this measurement, m and i can be defined as follows:

- $m = T_{pre}/s$ where T_{pre} is the number of pre-IDS transactions and s is the elapsed processing time (in seconds) for the transaction set
- $i = T_{post}/s$ where T_{post} is the number of post-IDS transactions and s is the elapsed processing time (in seconds) for the transaction set

Pre-IDS results show 408.38 as the volume of data transactions per second. Overall the system performed as expected with the greatest latency observed at iteration 4 with 309.67 as the volume of data transactions per second or a delta of 98.71 transactions.

$$v = m - i \quad (3)$$

4.4. Adaptability Rate

The Adaptability (a) Rate measurement determines the rate by which the presented dIDS was able to identify new, valid (v) transactions and new, authorized users (u). It is presented, as an aggregate and identified as total new (n) transactions. Note that new information always carries a higher rate of false positives when introduced to the dIDS and this will be reflected in the development of this measurement. The anticipated adaptability rate was achieved at $\geq 25\%$ of new transactions identified.

The threshold values for the Adaptability Rate were manually modified in the code as the iterations pro-

gressed through the testing phase. The policy tables were also modified manually to test the response of the dIDS and its ability to recognize changes and permissions.

Transaction volume was not considered in the intrusion probability calculation in an effort to keep the intrusion process independent from the input. This allowed the system to recognize a normalized range of probabilities regardless of transaction volume during the intrusion detection process. As the metrics contained within this research demonstrate, the probability range remained stable and consistent throughout the testing phase.

Using the generated data, the Apriori Hybrid data mining Algorithm was deployed using the SPMF software created by [30]. The resulted in an in the generation of 13,199 patterns of data. Examination of these patterns reflect a support of < 0.8 and creating less meaningful information. Often patterns were generated that had no value in the intrusion detection process. To avoid this, the threshold of pattern utilization was set at ≥ 0.8 . The results were 14 common behavior patterns and 95 data patterns. Within the data patterns 14 common fields were identified. These results were then utilized as the baseline signature of the dIDS system.

To support the incoming transaction set, 925 incoming transactions were generated using the Data Generator software. However, the study only used 115% or 12.5% of the amount of input that [7] used, the sample reduction did not affect the testing since the study is focused on the percentage of transactions identified as an intrusion and not the volume. Additionally, the following results can further be inferred to account for the larger dataset.

The data was generated based on the behaviors identified in the Apriori Hybrid data mining process. The patterns mined were determined to be the valid signatures while the remaining. Using the valid signatures patterns as the catalyst, both valid and invalid transactions were generated.

Monitoring and subsequent analysis of incoming transactions will determine probability of intrusion by identifying the tokens in a transaction. Tokens are included the action SELECT, INSERT, UPDATE, DELETE as well as the individual columns, with the identity of the requesting entity and WHERE and SET clauses, (when present). Once identified, these tokens were then compared to policy tables to determine if they were contained in an existing policy. If the transaction is validated based on the information in the policy tables, the transaction was considered a non-intrusion. If not validated, the intrusion process will be triggered.

Once triggered, the intrusion process will compare the incoming transaction to the database transaction logs in an attempt to identify a new, valid signature. The method of identifying the information contained within the logs was implemented by way of a dynamically generated

SELECT statement where a probability was computed based on how many of the tokens were found on a single log entry when taking into account the environmental factors (as previously noted). Next a probability was generated using the Stochastic Gradient Boosting algorithm. If the result was greater than 0.75, the transaction was marked as a new signature; otherwise it was subjected to further scrutiny. The 0.75 probability threshold used to identify a new signature was established at a high value since a transaction that is not yet considered valid can hold a higher degree of risk.

If the transaction was not validated in the new signature identification process, the transaction was subjected to the intrusion detection process. Unlike the new signature identification process, logs do not support the intrusion transaction. To refine the intrusion probability the transaction's probability of intrusion was computed by the Stochastic Gradient Boosting algorithm. This allowed control over the refinement of the threshold level. Then the transaction was subjected to further refinement using the Bayesian Belief Net algorithm. If the probability of this primary review was below the threshold, it was deemed as a non-intrusion event and added to the signature table as a validated signature. Should the incoming transaction require further analysis the same principle as identified for the stochastic Gradient Boosting Algorithm was applied to the application of the Bayesian Belief Net algorithm. If this probability computation resulted in a value that was above the control threshold, the transaction was deemed an intrusion and logged to the intrusion table to be used in the on-going signature identification process. There were six (6) iterations of the transaction cycle performed during the testing phase using the various threshold values.

$$a = (u + v) / n \quad (4)$$

5. Conclusions and Future Work

5.1. Conclusions

Based on the information provided to the dIDS model, the following conclusions have been drawn. Identification of "good" vs. "malicious" transactions is greatly dependent upon the information contained within the transaction itself, the log information, the number of rows being impacted by the database request and the policies pertaining to the entity making the request and the behavior of the requesting entity. Without consideration of each of these components within the context of each other, it cannot be accurately determined if the authorized insider is behaving as expected.

When using the Stochastic Gradient Boosting algorithm alone, when compared to using both algorithms to reach a finer degree of analysis, is less effective than

use both algorithms together in a single process when attempting to identify an intrusive transaction. By using both algorithms within the same intrusion process, the number of false-positives was markedly reduced.

Several types of transactions were introduced to indicate a new entity; the system was able to identify the new entity as the testing iterations progressed regardless of what the threshold rate was set at. When policy changes were introduced, the system correctly identified the intrusion and non-intrusion state.

Overall, the system presented proved very successful. Each goal with the exception of the maintaining the latency factor at a steady rate was met. Further research in the area of maintaining or reducing the latency of a Database intrusion detection system is warranted.

As has been observed, to use one methodology in an attempt to identify the insider threat phenomena in the context of the database environment, that supports a reasonable probability measurement, cannot be considered a complete solution. The uncertainty of the requestor's prior behavior must take into consideration along with the complete set of data and environmental factors in order to reach the conclusion that the insider is behaving beyond the boundaries as stated within defined security policies.

This research also observed that it is possible to leave the system in unattended learning environment in order to determine the probability of intrusion when the system is presented with new information as long as the other factors as noted are considered.

5.2. Future Works

The success of this research was based upon the research of many other researchers in not only intrusion detection but in the database technologies as well.

This research simply laid the foundations for future work to be investigated with respect to Database intrusion detection systems. To further the positive results presented, it can be expected that research in the following areas will build upon what has been presented in the preceding sections.

The presented research was based upon current research in intrusion detection models. Some of these models are utilized at the network level while others are at the database level. While intrusion detection does aid in the discovery of potential intrusions, it still requires a manual decision to be made by an intrusion administrator as to whether the incoming database request can be definitively considered a non-intrusion event. This is most often accomplished by some form of human intervention. Expansion of this study to move from an intrusion detection model to an intrusion prevention model allows for an expansion in the research area which enables the next

- for Association Rule Mining—A General Survey and Comparison,” *ACM SIGKDD Explorations Newsletter*, Vol. 2 No. 1, 2000, pp. 58-64.
[doi:10.1145/360402.360421](https://doi.org/10.1145/360402.360421)
- [25] P. H. Sharrod, “TreeBoost: Stochastic Gradient Boosting,” 2003. <http://www.dtreg.com/treeboost.htm>
- [26] P. J. Windley “Digital identity,” O’Reilly, Sebastopol, 2005.
- [27] S. Axelsson, “Combining a Bayesian Classifier with Visualization: Understanding the IDS,” *Proceedings of ACM Workshop on Visualization and Data Mining for Computer Security*, New York, 2004, pp. 99-108.
- [28] A. H. R. Karim, R. M. Rajatheva and K. M. Ahmed, “An Efficient Collaborative Intrusion Detection System for MANET Using Bayesian Approach,” *Proceedings of 9th ACM International Symposium on Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM 06)*, New York, 2006, pp. 187-190.
- [29] P. Mell, V. Hu, R. Lippman, J. Haines and M. Zissman, “An Overview of Issues in Testing Intrusion Detection Systems” 2003.
<http://csrc.nist.gov/publications/PubsNISTIRs.html>
- [30] P. Fournier-Viger, “Computer Software Documentation,” 2008. <http://www.philippe-fournierviger.com/spmf/>
- [31] United States of America (USA), “US Government Protection Profile: Intrusion Detection System for Basic Robustness Environments,” National Security Agency (NSA), Washington DC, 2007.