Admonita: A Recommendation-based Trust Model for Dynamic Data Integrity

Wassnaa Al-Mawee[1], Steve Carr[1] and Jean Mayo[2]
[1]*Department of Computer Science, Western Michigan University, 1903 W. Michigan Ave., Kalamazoo, MI 49008-5466, U.S.A.*
[2]*Department of Computer Science, Michigan Technological University, 1400 Townsend Dr., Hougton, MI 49931-1292, U.S.A.*

Keywords: Data Integrity, Biba Model, Trust, Subjective Logic, Tranquility.

Abstract: Data integrity is critical to the secure operation of a computer system. Applications need to know that the data that they access is trustworthy. Many current production-level integrity models are tightly coupled to a specific domain, (*e.g.*, databases), or only apply after the fact (*e.g.*, backups). In this paper we propose a recommendation-based trust model, called Admonita, for data integrity that is applicable to any structured data in a system and provides a measure of trust to applications on-the-fly. The proposed model is based on the Biba integrity model and utilizes the concept of an Integrity Verification Procedure (IVP) proposed by Clark-Wilson. Admonita incorporates subjective logic to maintain the trustworthiness of data and applications in a system. To prevent critical applications from losing trust, Admonita also incorporates the principle of weak tranquility to ensure that highly trusted applications can maintain their trust levels. We develop a simple algebra around these elements and describe how it can be used to calculate the trustworthiness of system entities. By applying subjective logic, we build a powerful, artificial and reasoning trust model for implementing data integrity.

1 INTRODUCTION

Computer security is at the forefront of modern society. The Ponemon Institute and IBM Security have published their *2019 Cost of a Data Breach Report* showing an average cost of almost $4 million. With this cost, organizations must do everything they can to ensure the confidentiality, availability and integrity of data in their systems. In order to do this, researchers must develop practical models that allow a system to provide assurance of data security.

Confidentiality and availability have justly received a significant amount of attention in the literature. However, these two factors are not enough. For instance, encryption and access control make it difficult for unauthorized access to data, but do not verify the integrity of the data being protected. An authorized user or application may make changes to data that lower its integrity and confidentiality measures are unable to detect the authorized, yet erroneous change. Thus, an application may have access to data that may have been altered in a way that lowers the integrity of the data without being able to detect the problem.

Integrity often is maintained by restricting access to high integrity items to only subjects that have high integrity. However, as illustrated in the previous paragraph, integrity is also a property of the data itself, not just of who accesses or modifies it. How can an application know how trustworthy the data it accesses is? In addition, if an application tries to access data that is not very trustworthy, should that application be allowed to access that data and, if so, does the access affect the future trustworthiness of both the data and subject?

This paper addresses these questions by using an improvement to the trust-enhanced data integrity model of Oleshchuk (Oleshchuk, 2012). We present a recommendation-based trust model for dynamic data integrity, called Admonita. Admonita is based upon subjective logic (Jøsang, 1999; Jøsang, 2001; Jøsang, 2002; Gao et al., 2009; Oleshchuk, 2012) and the Biba integrity model (Biba, 1977); however, it incorporates the idea of an Integrity Verification Procedure (IVP) from the Clark-Wilson model (Clark and Wilson, 1987), the principle of tranquility (Bell and Padula, 1973; Bishop, 2019) that allows integrity levels to increase or decrease, and the notion that the data

Al-Mawee, W., Carr, S. and Mayo, J.
Admonita: A Recommendation-based Trust Model for Dynamic Data Integrity.
DOI: 10.5220/0010150402730282
In *Proceedings of the 7th International Conference on Information Systems Security and Privacy (ICISSP 2021)*, pages 273-282
ISBN: 978-989-758-491-6

itself has a measure of integrity apart from who modifies it.

In Admonita, the trust level for subjects and objects is set by a trusted authority. Admonita then incorporates the opinion of an independent observer via an IVP implemented in a language that describes what it means for structured data to have integrity (Bonamy, 2016; **?**). Admonita maintains the trustworthiness of both subjects and objects in a computing system via the conjunctive, consensus and recommendation operators from subjective logic. Admonita adjusts the trustworthiness of entities dynamically based upon the trust levels of subjects and the objects they access, and includes bidirectional weak tranquility to allow the trust levels to increase or decrease.

The rest of paper is organized as follows. In Section 2 we give the background related to Biba, Clark-Wilson, and data integrity language. In Section 3 we give an overview of related work on trust models based subjective logic. In Section 4 we formally introduce needed notation, relations and notions of subjective logic. Then we present the tranquility principle for dynamic data protection in Section 5. A description of our proposed recommendation-based trust model is presented in Section 6. Section 7 provides a structure and example for trust authentication in our trust model. Finally, Section 8 concludes this paper.

2 BACKGROUND

2.1 Security Models

There are several security models that address integrity for secure systems. The most directly useful and related to our work are Biba and Clark-Wilson. Each integrity model offers a definition of data integrity and introduces their own mechanisms for preserving integrity.

The first model that supported data integrity based on a subject's static integrity level was the Biba model developed by Kenneth J. Biba in 1977 (Balon and Thabet, ; Biba, 1977). The model describes a set of subjects, a set of objects, and a set of integrity levels. Subjects may be either users or processes. Each subject and object is assigned an integrity level, denoted as $I(S)$ and $I(O)$, for the subject S and the object O, respectively. The idea is that subjects with lower integrity levels are not permitted to modify objects that have higher integrity levels. Similarly, subjects with high integrity levels cannot be corrupted by objects with low integrity levels. Biba is a well-known general integrity model in computer systems. Its mandatory integrity property succeeds at enforcing integrity in a system, but it does not deal with the integrity of data itself; authorized users can still make improper modifications. For example, if a trusted user account is compromised, an attacker can use a trusted user's integrity level to modify high-level integrity resources.

One of the biggest threats to a company's data are its employees, including users and administrators. They are able to access data, make modification and copies, use USB discs etc. A popular adoption of Biba model is in modern Microsoft Windows operating systems where processes carry integrity labels and low-integrity subjects/processes cannot interact with high-integrity ones. Windows mandatory access control (MAC) ensures data integrity via an access control mechanism. Windows restricts access rights depending on whether the subject's integrity level is equal to, higher than, or lower than the object's integrity level. The integrity level of an object is stored as a mandatory label access control entry (ACE) that distinguishes it from the discretionary ACEs governing access to the object (Microsoft,). The limitation of this technology is that the ACEs can be modified by an offline attack (modification by the system's administrator). This problem can be solved using Admonita. Admonita is recommendation-based model that uses past behavior to determine whether to trust an entity. The independent trust opinion of a declarative system that states what it means for data to have integrity results from validating the actual user input against the resources that he/she wants to access. The trust level is independently computed without human intervention and combined with the integrity levels of subjects and objects using subjective logic.

A second integrity model is that proposed by David Clark and David Wilson (CWM) (Clark and Wilson, 1987). CWM focuses on the prevention and detection of data integrity faults using transactions. The model is based on two concepts that are used to enforce commercial security polices or constraints:

1. Well-formed transactions: A user can manipulate data using only constrained rules to ensure the integrity of data.
2. Separation of duty among users: A person who has the permission to perform well-formed transactions may not have the permission to access the constrained data.

CWM enforces integrity controls on data by separating all the data items within a system into two groups:

1. Constrained Data Items (CDIs): Data items that have associated integrity constraints.

2. Unconstrained Data Items (UDIs): Data items that do not have associated integrity constraints.

After classifying the data items, the integrity system tests the data items through two types of procedures:

1. Transformation Procedures (TPs): Achieve data transactions by changing the system's CDIs from one valid state to another.
2. Integrity Verification Procedures (IVPs): Ensure that all the CDIs conform to the integrity constraints or specifications.

CWM provides data integrity but imposes a number of restrictions that make it impractical to implement. A transformation procedure may have an issue if a single application is able to execute many different transformations. For example, a text editor can be used to produce HTML files, or to edit the UNIX password file. To implement CWM, the text editor must be broken into an HTML editor, and a password file editor to be certified to produce valid HTML files, and valid UNIX password file. Additionally, an administrator needs to manage and verify all the editors and that is impractical.

With respect to integrity protection, all data integrity models deal with the preservation of trust. There is a need for a more flexible definition of data integrity that takes into account whether the data itself can be trusted apart from who modifies the data.

2.2 A Data Integrity Language

We have incorporated a data integrity language, which we will call Maia (Bonamy, 2016; **?**) to work as an IVP. A Maia specification is compiled into an authorized program that ensures that all the constrained data conform to the integrity constraints or specifications contained in the Maia specification file. Maia is a specification language that declares what it means for arbitrary structured file types (Bonamy, 2016; **?**) to have integrity as a property of the data contained within the file itself. For example, Maia within the context of Linux has been used to specify the integrity of system configuration files, PNG files, and others.

In Maia, a file verification process is accomplished using two phases that correspond to checking the file syntax and semantics. In the first phase, the user provides a grammar in order to verify the file structure and extract its syntactic elements for processing. This syntax-checking component of Maia is designed to work like a normal parser as generated by a parser generator where the elements of the file are put into collections. The second phase of Maia checks the collections of data in the syntactic elements by using set theory and predicate calculus to express the integrity constraints.

3 RELATED WORK

The following are different enhanced trust-based subjective logic models to support various organizational security policies that have been proposed. Oleshchuck proposes a trust-enhanced data integrity model that is based on the Biba integrity model using subjective logic (Oleshchuk, 2012). In his model, he reformulates the rules of the Biba integrity model in terms of trust and proposes how to combine Role-Based Access Control RBAC with the introduced integrity model. Gao, et al. (Gao et al., 2009), propose a trust model by analyzing and improving subjective logic. By using subjective logic in their model (Jøsang, 2002), they can evaluate the trust relationship between peers and resolve security problems in practical computing environments. Jøsang proposes a trust management system based on subjective logic (Jøsang, 2001). He proposes an evidence space and opinion space that are used to evaluate and measure trust relationships. These policy-based trust models use credentials to instantiate policy rules that determine whether to trust an entity, resource or information. The policies do not protect the system entities since the credentials themselves are information that is not protected by the model. On the other hand, trust models preserve the initial evaluation of data integrity by providing information about the trustworthiness of data and entities. These models do not consider all of the side effects of dynamic data integrity. For instance, the trust opinions of the system's subjects can keep obtaining lower trust levels when they read less trusted data but there is not mechanism in the model to raise the integrity levels, possibly resulting in isolation of the subject.

To deal with the problems mentioned above, we propose a new recommendation-based trust model that is based on subjective logic and bi-directional weak tranquility. The recommendation-based trust model uses past behavior during interactions and information from other resources to determine whether to trust an entity. Our model adopts the rules from the Biba integrity model and incorporates recommendation opinions from Maia. In our trust model, integrity levels of subjects and objects are expressed as trust opinions. Since the security of the system is a subjective measure that depends on individuals who are qualified to express trust opinions, we define such trust opinions in the framework of subjective logic. We compute the recommendation values and the trust

opinions with a simple algebra based on the trust metrics of our model. Also, we add a flexible definition of data integrity by using bidirectional weak tranquility.

4 SUBJECTIVE LOGIC

In this section, we use an artificial reasoning framework called subjective logic to express the levels of trust. Due to the lack of certainty about the degree of the trustworthiness of subjects and objects, we need to have opinions to measure the integrity of these subjects and objects. Subjective logic defines the term *opinion*, w, which expresses an opinion about the trust level of subjects/objects (Jøsang, 1999; Jøsang, 2002). The opinion translates into degrees of trust, distrust as well as uncertainty, that represents the absence of both trust and distrust values. Let t, d, and u be *trust*, *distrust* and *uncertainty*, respectively, such that:

$$t+d+u=1, \quad t,d,u \in [0,1] \tag{1}$$

The opinion $w=\{t,d,u\}$ is a triplet satisfying (1). We use opinions to express trust levels. Having different levels of trust instead of a single level, such as in Biba, provides a better integrity model for real-world applications.

Subjective logic defines set logical operators that are equivalent to traditional logical operators, such as conjunction (AND), disjunction (OR), and negation (NOT), as well as some non-traditional operators that are used for combining opinions, such as recommendation and consensus. The expressed opinions are the input and output parameters for subjective logic operators. For the purpose of this paper, we will define only consensus, recommendation and conjunctive operators.

Let A and B be two entities that represent observers who maintain the trust opinions of system resources, and let o be an object. When there are independent opinions about o, subjective logic defines a *consensus* operator to combine these independent opinions.

Let $w_o^A=\{t_o^A,d_o^A,u_o^A\}$ and $w_o^B=\{t_o^B,d_o^B,u_o^B\}$ be opinions held by the observers A and B, respectively, about o. According to subjective logic, the combined consensus opinion $w_o^{A,B}$ based on opinions w_o^A and w_o^B is defined as follows:

$$\begin{gathered} w_o^{A,B}=w_o^A\oplus w_o^B \\ =\{t_o^{A,B},d_o^{A,B},u_o^{A,B}\} \\ where, \\ \left\{\begin{array}{c} t_o^{A,B}=(t_o^Au_o^B+t_o^Bu_o^A)/(u_o^A+u_o^B-u_o^Au_o^B) \\ d_o^{A,B}=(d_o^Au_o^B+d_o^Bu_o^A)/(u_o^A+u_o^B-u_o^Au_o^B), \\ u_o^{A,B}=(u_o^Au_o^B)/(u_o^A+u_o^B-u_o^Au_o^B) \end{array}\right\} \end{gathered} \tag{2}$$

Let A and B be two observers such that observer A invokes observer B to access an object o. Let $w_B^A=\{t_B^A,d_B^A,u_B^A\}$ be A's opinion about B's recommendation, and let $w_o^B=\{t_o^B,d_o^B,u_o^B\}$ be B's opinion about the trustworthiness of the object o. Subjective logic defines a *recommendation* operator to compute the indirect opinion w_o^{AB} based on opinions w_B^A and w_o^B as:

$$\begin{gathered} w_o^{AB}=w_B^A\otimes w_o^B \\ =\{t_o^{AB},d_o^{AB},u_o^{AB}\} \\ where, \\ \left\{\begin{array}{c} t_o^{AB}=(t_B^At_o^B) \\ d_o^{AB}=(t_B^Ad_o^B), \\ u_o^{AB}=(d_B^A+u_B^A+t_B^Au_o^B) \end{array}\right\} \end{gathered} \tag{3}$$

Furthermore, subjective logic defines the *conjunctive* operator that expresses an opinion that is held by observer A about the trustworthiness of two distinct objects o_1 and o_2. Let $w_{o_1}^A=\{t_{o_1}^A,d_{o_1}^A,u_{o_1}^A\}$ and $w_{o_2}^A=\{t_{o_2}^A,d_{o_2}^A,u_{o_2}^A\}$ be observer A's opinions about o_1 and o_2. Then the conjunction opinion $w_{o_1\wedge o_2}^A$ of $w_{o_1}^A$ and $w_{o_2}^A$ is defined by:

$$\begin{gathered} w_{o_1\wedge o_2}^A=w_{o_1}^A\wedge w_{o_2}^A \\ =\{t_{o_1\wedge o_2}^A,d_{o_1\wedge o_2}^A,u_{o_1\wedge o_2}^A\} \\ where, \\ \left\{\begin{array}{c} t_{o_1\wedge o_2}^A=(t_{o_1}^At_{o_2}^A) \\ d_{o_1\wedge o_2}^A=(d_{o_1}^A+d_{o_2}^A-d_{o_1}^Ad_{o_2}^A), \\ u_{o_1\wedge o_2}^A=(t_{o_1}^Au_{o_2}^A+u_{o_1}^At_{o_2}^A+u_{o_1}^Au_{o_2}^A) \end{array}\right\} \end{gathered} \tag{4}$$

In our model, we consider the opinion w_B, where $d_B<t_B$, to be more trustworthy than opinion w_A , where $d_A<t_A$, denoted $w_B\gg w_A$, if and only if $t_B>t_A$. When $t_B=t_A$, we consider higher uncertainty to be more trustworthy. In the case of $t_B=t_A \implies w_B\gg w_A \iff u_B>u_A$.

5 TRANQUILITY FOR DYNAMIC INTEGRITY POLICY

In this section, we outline a tranquility principle that is trust-enhanced to protect trust levels of system entities and resources.

When a model, such as that of Oleshchuk, allows the trustworthiness of subjects to decrease due to reading low trusted data, the subject may become isolated from system resources since Biba's model incorporates only unidirectional weak tranquility. Such isolation can cause a violation of the security policy. For example, the *ls* Linux command is a command-line utility for listing the contents of a directory or directories given to it via standard input, and it writes

to the standard output (William E. Shotts,). When *ls* accesses a corrupted directory/file, the trust level of *ls* after will decrease. If *ls* continues accessing objects with low integrity levels, the trust level of *ls* may become isolated from system resources. To solve this issue, we apply the principle of weak tranquility such that the trust level may both increase and decrease, making it bidirectional.

The tranquility principle allows controlled copying from high security levels to low security levels via trusted subjects. There are two forms of the tranquility principle: *strong tranquility* and *weak tranquility*. In strong tranquility, the security levels do not change during the normal operation of the system. In weak tranquility, the security levels may never change in such a way as to violate a defined security policy. Bidirectional weak tranquility is more desirable in our model. An entity may obtain a new low trust level due to accessing low integrity data or invoking low integrity entities. By applying bidirectional weak tranquilly, the entity can progressively accumulate higher trust levels, as actions require it. In other words, subjects and objects integrity levels will be managed within an allowable range to make the process more flexible in application. So, our model not only incorporates weak tranquility in a bi-directional manner, the are both maximum and minimum trusts levels that represent boundaries across which an object's integrity level may not change.

6 RECOMMENDATION-BASED TRUST MODEL FOR DATA INTEGRITY

Trust models are divided into two types: policy-based models and recommendation-based models. Both types use a language to express relationships about trust. Each type provides a measure of the trust in an entity, and the result of the evaluation is a complete *trust*, a complete *distrust*, or somewhere between *certain* or *uncertain*.

Policy-based models require a language in which to express and analyze system policies. For example, the Keynote trust management system (Blaze et al., 1998) that is based on Policy-Marker (Blaze et al., 1996) is extended to support applications that use public keys. Recommendation-based models use past behavior to determine whether to trust an entity, including recommendations from other entities. For example, Abdul-Rahman and Hailes (Abdul-Rahman and Hailes, 1997) base trust on the recommendations of other entities. In their model, they consider direct trust relationships and recommender trust relationships. Trust is computed based on integer values. They use -1 for direct trust as representing untrusted, values from 1 to 4 as representing the lowest to highest trust values, and 0 as the inability to make trust judgments. For recommender trust values, the integers -1 and 0 have the same meaning as with direct trust, while the values from 1 to 4 indicates how close the recommender judgment is to the entity that is being recommended.

Admonita is a recommendation-based trust model. It is based on Biba and Maia. In our proposed model, the Biba integrity model defines the subject-objects access properties, while Maia works as an Integrity Verification Procedure IVP that preserves data integrity. Basically, a Maia specification defines a set of constraints declaring what it means for data to have integrity. Maia verifies structured data when a subject writes to the file and generates a limited number of integrity levels to reflect the evaluation of the data's integrity.

The Biba integrity model is concerned with an unauthorized modification of data within a system by controlling who may access it. It works as a prevention system for data integrity. The model deals with a set of subjects, a set of objects, and a set of integrity levels. Subjects may be either users or processes. Each subject and object is assigned an integrity level, denoted as $I(s)$ and $I(o)$, for the subject s and the object o, respectively. The integrity levels describe how subjects and objects are more or less trustworthy regarding a higher or lower integrity level.

Let $S = \{s_1, s_2, \dots\}$ be a set of subjects, and $O = \{o_1, o_2, \dots\}$ be a set of objects. According to subjective logic, the opinions about a subject and an object are expressed as $w_s = \{t_s, d_s, u_s\}$ and $w_o = \{t_o, d_o, u_o\}$ respectively, where $s \in S$ and $o \in O$. Therefore, the trust opinion about the subject w_s represents the integrity of the subject $I(s)$. Similarly, the trust opinion about the object w_o represents the integrity of the object $I(o)$.

According to (Gambette, 1988), the definition of trust is "Anna trusts Bernard if Anna believes, with the level of subjective probability, that Bernard will perform a particular action, both before the action can be monitored (or independently of capacity of being able to monitor it) and in a context in which it affects Anna's own action." If Anna establishes trust in Bernard based on her observation and other interactions, the trust is *direct*. If it is established based on Anna's acceptance of Bernard's recommendation of other entities, then the trust is *indirect*.

Admonita combines *direct* and *indirect* opinions about the trustworthiness of subjects and objects. A

security officer T expresses direct trust opinions about the subject s, denoted as w_s^T , and the trust opinions about the object o, denoted as w_o^T . Also, T maintains a list of minimum trust opinions for subjects, denoted as w_{s-min}^T, and list of maximum trust opinions for objects, denoted as w_{o-max}^T. Maia expresses the indirect subject-object trust opinion, denoted as $w_{s_o}^M$.

We incorporate the Biba model operations for both subjects and objects:

- Observe: Allows a subject s to read information in an object o, denoted as $read(s,o)$.
- Update: Allows a subject s to write or update information in an object o, denoted as $update(s,o)$.
- Invoke: Allows a subject s_1 to execute another subject s_2, denoted as $invoke(s_1,s_2,o)$.

The Biba model can be divided into two types of policies, mandatory and discretionary. Most literature on the Biba model refers to the model as being mandatory as a part of the strict integrity policy (Balon and Thabet,). The Biba model defines a number of rules as part of the strict integrity policy. We reformulate each rule and compute the integrity level of subjects $I(s)$ and objects $I(o)$ as follows:

6.1 Simple Integrity Property

The Simple Integrity Property enforces *no-read-down*. It allows a subject to read (observe) an object only if the integrity level of the subject is less than the integrity level of the object.

$$s \in S \;\; reads \;\; o \in O \iff I(s) \leqslant I(o)$$

This ensures that high-integrity data cannot be directly contaminated by low-integrity data. For example, if the simple integrity rule is enforced, a low-integrity process may read high-integrity data, but it cannot contaminate itself by reading low-integrity data.

Our trust model adds a dynamic property to the Simple Integrity Property to allow high trust subjects to access low trust objects, however, the integrity of the subject may be lowered. The Simple Integrity Property in our model is reformulated as described below.

$$\begin{aligned} &\forall s \in S, \forall o \in O: \;\; read(s,o) \iff \\ &if \;\; I(s) > I(o) \;\; then \;\; I'(s) = I(s) \otimes I(o) \end{aligned} \tag{5}$$

When s reads o, denoted $read(s,o)$, the integrity of s may be changed by o, while the integrity of o will not be changed. If s reads less trusted data, then the integrity level of s after reading, denoted $I'(s)$, will decrease. To compute the indirect opinion $I'(s)$, denoted as $w_{s_o}^{TM}$, let T and Maia be two observers such that the observer T invokes observer Maia to access an object o. Let w_M^T be T's opinion about Maia's recommendation, denoted as $w_{s\wedge o}^T$, and let $w_{s_o}^M$ be Maia's opinion about the trustworthiness of the object o that is accessed by the subject s. We adjust reasoning from (3) and (4) and we argue the applicability of *conjunction* and *recommendation* operators described in the previous section as follows:

$$\begin{aligned} w_{s_o}^{TM} &= w_{s\wedge o}^T \otimes w_{s_o}^M \\ &= (w_s^T \wedge w_o^T) \otimes w_{s_o}^M \end{aligned} \tag{6}$$

T expresses the direct trust opinions of w_s^T and w_o^T. These two opinions are combined using the conjunctive operator since they are both are assigned by the same observer. $w_{s\wedge o}^T$ represents T's opinion about Maia's recommendation. On the other hand, the Maia model expresses the indirect trust opinion about $w_{s_o}^M$.

To compute the indirect opinion about the trustworthiness of $w_{s_o}^{TM}$ based on the trustworthiness of recommendation of Maia, the two opinions are combined using the *recommendation* operator. Equation (6) will decrease the integrity level of s when it reads less trusted o.

To prevent subject isolation, bidirectional weak tranquility is applied to ensure that the new obtained trust level is within an allowable range. It is accomplished by comparing the new trust value $I'(s)$, denoted as t_s', against its minimum value of trust, denoted as t_{s-min} as follows:

1. If the new trust value of the subject is greater than its minimum value such that $t_s' > t_{s-min}$, then read access will be granted.
2. If the trust values of a subject are equal such that $t_s' = t_{s-min}$, we consider higher uncertainty to be more trustworthy. In the case when $t_s' = t_{s-min}$ then $w_s' > w_{s-min}$ if $u_s' > u_{s-min}$.
3. If the new trust value of the subject is less than its minimum value such that $t_s' < t_{s-min}$ then the new integrity level of the subject violates the integrity policy. To solve this problem, we consider two cases:
   - If the subject is not allowed to have an integrity level less than its minimum integrity level, then read access will be denied and that reflects the *no-read-down* rule of the Biba integrity model.
   - If the subject is allowed to have an integrity level less than its minimum integrity level then the read access will be granted and the integrity level of the subject will be forced to go back to its previous integrity level to prevent subject isolation.

$$if \;\; t_s' < t_{s-min} \;\; then \;\; I'(s) = I(s) \tag{7}$$

6.2 Integrity Star Property

The second property of a Biba policy enforces *no-write-up*. It allows a subject to write an object only if the integrity level of the object is less than or equal to the integrity level of the subject.

$$s \in S \;\; updates \;\; o \in O \iff I(o) \leqslant I(s)$$

The Integrity Star Property in our model can be reformulated as follows:

$$\begin{array}{c} \forall s \in S, \forall o \in O: \;\; update(s,o) \iff \\ if \;\; I(o) \leqslant I(s) \;\; then \;\; I'(o) = I(o) \oplus I(s) \end{array} \tag{8}$$

When s updates o, denoted $update(s,o)$, the integrity level of o, denoted as $I(o)$, with trust opinion w_o will be changed by the integrity level of s, denoted as $I(s)$, with trust opinion w_s. If w_s is more trustworthy than w_o then the integrity level of the object after the update, denoted as $I'(o)$, will be increased and $I(s)$ will not change. In contrast, if w_s is less trustworthy than w_o then s will not be allowed to update o. This corresponds to the *no-write-up* rule of the Biba integrity model.

We consider two scenarios to enforce the *no-write-up* rule. First, when a high trust subject s updates a low trust object o invalidly, either accidentally or intentionally, Maia generates a lower recommendation opinion with a higher *distrust* value for accessing the object o. That lowers the integrity level of s by updating $I'(s)$ using (6). Then, $I'(s)$ is compared against $I(o)$ without updating $I'(o)$ with (8). With that, our model enforces *no-write-up*, and s will be denied to update o. After that, $I'(s)$ will be set using (7), to avoid isolation from the system resources.

The second scenario, occurs when a high trust subject s updates a low trust object in a valid format. In this case, the Maia generates a valid recommendation opinion for accessing the object o.

$$\begin{array}{l} w_{s_o}^{T,M} = w_{o \wedge s}^{T} \oplus w_{s_o}^{M} \\ = (w_o^T \wedge w_s^T) \oplus w_{s_o}^{M} \end{array} \tag{9}$$

T expresses the direct trust opinions of w_o^T and w_s^T. These two opinions are combined using the conjunctive operator since they are both are assigned by the same observer. $w_{o \wedge s}^T$ represents T's opinion about Maia's recommendation. As with Simple Integrity Property, Maia expresses the indirect recommendation trust opinion about $w_{s_o}^M$. Since the Integrity Star Property in our model keeps s unchanged, we introduce a *conjunctive consensus* term $w_{s_o}^{T,M}$ that combines two independent opinions about accessing o. Equation (9) enforces increasing the integrity level of o when it is accessed by highly trustworthy s.

To prevent object isolation, bidirectional weak tranquility is applied to ensure that the new obtained trust level $I'(o)$ is within an allowable range. It is accomplished by comparing the new trust value of $I'(o)$, denoted as t'_o, against its maximum value of trust, denoted as t_{o-max}, as follows:

1. If the new trust value of the object is less than its maximum value such that $t'_o < t_{o-max}$ then update will be granted.
2. If the trust values of objects are equal such that $t'_o = t_{o-max}$, we consider higher uncertainty to be more trustworthy. In the case of $t'_o = t_{o-max}$ then $w'_o > w_{o-max}$ if $u'_o > u_{o-max}$.
3. If the new trust value of the object is greater than its maximum value such that $t'_o > t_{o-max}$ then the new integrity level of the object violates the integrity policy. To solve this problem, we consider two cases:
   - If the object is not allowed to have an integrity level greater than its maximum integrity level, then the update will be denied.
   - If the object is allowed to have an integrity level greater than its maximum integrity level, then the update access is granted and the integrity level of the object will be forced to go back to its previous integrity level to prevent object isolation.

$$if \;\; t'_o > t_{o-max} \;\; then \;\; I'(o) = I(o) \tag{10}$$

6.3 Invocation Property

In Biba's model, a subject may execute another subject at its own integrity level or below.

$$s_1 \in S \;\; invokes \;\; s_2 \in S \iff I(s_2) \leqslant I(s_1)$$

This last property states that a subject at one integrity level is prohibited from invoking (send/request messages for service) a subject at a higher level of integrity. The Invocation Property in our model is reformulated as follows:

$$\begin{array}{c} \forall s_1, s_2 \in S, \forall o \in O: \;\; read(s_1, s_2, o) \iff \\ if \;\; I(s_1) < I(s_2) \;\; then \;\; I'(s_1) = I(s_1) \otimes (I(s_2) \otimes I(o)) \end{array} \tag{11}$$

In our trust model, when s_1 invokes s_2 to access o, it is denoted as $invoke(s_1, s_2, o)$. According to the strict integrity policy, our trust model requires $I(s_1) \geqslant I(s_2)$ preventing a less trustworthy s_1 using more trustworthy s_2 to update the data. This condition keeps $I(s_1)$ unchanged when s_1 reads lower integrity data o via s_2.

However, when $I(s_1) < I(s_2)$, the indirect opinion of $I'(s_1)$, denoted as $w_{s_{1_{s_{2_o}}}}^{TM}$, can be calculated using the conjunctive recommendation term as follows:

$$\begin{array}{l} w_{s_{1_{s_{2_o}}}}^{TM} = w_{s_1}^{T} \otimes (w_{s_2 \wedge o}^{T} \otimes w_{s_{2_o}}^{M}) \\ = w_{s_1}^{T} \otimes ((w_{s_2}^{T} \wedge w_o^T) \otimes w_{s_{2_o}}^{M}) \end{array} \tag{12}$$

In Equation (12), the security officer T expresses the direct trust opinions of $w^T_{s_1}$, $w^T_{s_2}$ and w^T_o. The opinions of $w^T_{s_2}$ and w^T_o are combined using the conjunctive operator since they are both assigned by the same observer. Maia expresses the indirect trust opinion $w^M_{s_{2o}}$. Then the indirect recommendation opinion of $w^{TM}_{s_{2o}}$ is combined with $w^T_{s_1}$ using the recommendation operator. As a result, the integrity level of s_1 decreases due to using highly trusted resources.

The proposed trust opinion calculations associated with *no-read-down* and *no-write-up* operations along with a dynamic range of integrity levels give our model more flexibility and better control of integrity violations.

7 A STRUCTURE AND EXAMPLE FOR TRUSTWORTHINESS AUTHENTICATION

Fig. 1 illustrates a possible structure for computing the integrity level as trust opinions about subjects/objects. The structure above the dotted line represents the opinions of the security officer T about subjects and objects as stored in $T's$ private database. Also, T maintains a list of minimum trust opinions for each subject s, denoted w^T_{s-min}, and list of maximum trust opinions for each object o, denoted w^T_{o-max}.

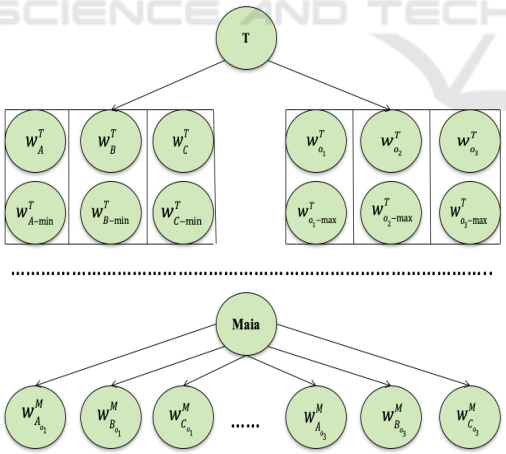

Figure 1: A Structure of Trustworthiness Authentication in Recommendation-Based Trust Model.

To ensure dynamic computation of trust opinions for system's entities, T assigns values of 0 and 1 for trusted subjects and objects, denoted T_s and T_o respectively. A value of 0 for a subject s does not allow s to obtain an integrity level less than its minimum integrity level. A value of 0 for an object o does not allow o to obtain an integrity level greater than its maximum integrity level. In contrast, a value of 1 for s allows s to have access to less trustworthy data and return to its previous integrity level. Similarly with o, a value of 1 allows o to have an integrity level greater than its maximum, for updates, and return to its previous integrity value.

The integrity level of a subject s reflects how much T trusts s when accessing an object. It is assumed that T knows the trust levels of s. On the other hand, the integrity level of an object reflects $T's$ opinion about the trustworthiness of the data itself.

In our model, T must keep a list of her opinions, w^T_s and w^T_o, about the trustworthiness of subjects and objects, respectively. T's opinions about an subject reflects the trust level of the subject. However, T's opinions w^T_{s-min} about s ensure dynamic data integrity while the T's opinion about an object reflects the trust level about the data itself. Table 1 gives an example of possible opinion values.

Table 1: Security Officer's Opinions about Subjects' Trustworthiness.

S	w^T_S	w^T_{S-min}	$T-S$
A	{1.00, 0.00, 0.00}	{0.99, 0.01, 0.00}	1
B	{0.98, 0.00, 0.02}	{0.85, 0.10, 0.05}	1
C	{0.88, 0.10, 0.02}	{0.80, 0.10, 0.10}	0

Table 2: Security Officer's Opinions about Objects' Trustworthiness.

O	w^T_O	w^T_{O-max}	$T-O$
o_1	{0.90, 0.05, 0.05}	{1.00, 0.00, 0.00}	0
o_2	{0.96, 0.02, 0.02}	{0.96, 0.02, 0.02}	0
o_3	{0.98, 0.00, 0.02}	{0.98, 0.00, 0.02}	1

Table 3: Maia's Opinions about (Subject-Object) Trustworthiness.

S_{o_1}	$w^M_{S_{o_1}}$
A_{o_1}	{0.95, 0.01, 0.04}
B_{o_1}	{1.00, 0.00, 0.00}
C_{o_1}	{0.89, 0.02, 0.09}

In order to enforce weak tranquility, T must maintain a list of her maximum opinions about objects w^T_{o-max}. Table 2 gives an example of possible opinion values. The structure below the dotted line represents a list of Maia trust recommendations $w^M_{S_o}$, based upon the Maia specification for that file, for each subject that wants to access the object o. Table 3 gives an example of possible opinion values.

Assume subject B wants to read an object o_1, $read(B, o_1)$. Since $I(B)$ is greater than $I(o_1)$, the trust of B, denoted $I'(B)$, can now be calculated using (6):

$$w^{TM}_{B_{o_1}} = (w^T_B \wedge w^T_{o_1}) \otimes w^M_{B_{o_1}} = \{0.882, 0.00, 0.118\}$$

Notice that the new integrity level of B, 0.882, is less than the old value 0.98 due to B reading object o_1 with a lower trust level than B. Also, the new $I'(B)$ satisfies (7) since it does not fall below its minimum trust value. Now, the integrity level of B in Table 1 will be replaced by the new value in order to prevent the low integrity of B from updating other objects in future interactions.

Suppose subject B wants to update object o_1, $update(B, o_1)$. Since $I(B)$ is greater than $I(o_1)$, the trust of o_1, denoted $I'(o_1)$, can now be calculated using (9):

$$w^{T,M}_{B_{o_1}} = (w^T_{o_1} \wedge w^T_B) \oplus w^M_{B_{o_1}} = \{1.00, 0.00, 0.00\}$$

Notice that the new integrity level of o_1, 1.00, is greater than the old value 0.90 since B has a higher trust value than o_1 . Also, $I'(o_1)$ satisfies (10) since it does not exceed the maximum trust opinion. Now, the integrity level of o_1 in Table 2 will be replaced by the new value.

Consider the case when a subject invokes another subject to access an object o. Assume B invokes A to access o_1. We need to modify the trust level of B for two reasons. First, the integrity level of B is lower than the integrity level of A, so the trust model will prevent B from using A. Second, subject A accesses less trusted data o_1. This decreases the integrity level of A.

To calculate the trustworthiness of B, $I'(B)$, first $I'(A)$ is calculated using (6) and (7) to let A obtain back its trust opinion since it is a trusted subject. Then $I'(B)$ is calculated using (12):

$$w^{TM}_{B_{A_{o_1}}} = w^R_B \otimes ((w^T_A \wedge w^T_{o_1}) \otimes w^M_{A_{o_1}}) = \{0.8379, 0.00882, 0.15328\}$$

The new integrity level of B, 0.8379, is less than the old value 0.98 and that is due to the Invocation Property. In addition, the new trust value of B violates the integrity policy since it is less than its minimum trusted value. However, B is a trusted subject. Therefore, our trust model allows B to read the less trusted object and obtain back its trust opinion. Now, the trustworthiness of B, denoted as $I'(B)$, can be calculated using (7):

$$if \;\; t'(B) < t_{B-min} \;\; then \;\; I'(B) = I(B)$$

$$I'(B) = \{0.98, 0.00, 0.02\}$$

If B is not a trusted subject, then the read access will be denied.

8 CONCLUSIONS AND FUTURE WORK

In this work we propose a new recommendation-based trust model for data integrity called Admonita. Admonita incorporates subjective logic, the Biba integrity model, the Clark-Wilson integrity model and the principle of bidirectional weak tranquility. Compared to previous models, our model adds the opinion of an IVP from Clark-Wilson of the integrity of the data that is a property of the data itself rather than the opinion of a trusted user. In addition, our model uses bidirectional weak tranquility to allow opinions about the integrity of data to change dynamically within a restricted range. The result is a model that determines the integrity of subjects and objects in a system that is not based solely on the integrity of the users in the system.

In the future, we plan to implement Admonita in a real system and measure its performance. This will involve creating a high-performance compiler for Maia that utilizes its natural parallelism. The result will be a system that measures and maintains the trust levels for the applications and data contained within it.

REFERENCES

Abdul-Rahman, A. and Hailes, S. ((1997). A distributed trust model. In *Proceeding of the 1997 Workshop on New Security Paradigms*, pages 48 – 60.

Al-Mawee, W., Carr, S., Bonamy, P., and Mayo, J. (2019). Maia: A language for mandatory integrity controls of structured data. In *Proceedings of the 5th International Conference on Information Systems Security and Privacy (ICISSP 2019)*.

Balon, N. and Thabet, I. The biba security model,. https://pdfs.semanticscholar.org/7360/c680906617622f27ef2596c7efcc902795db.pdf.

Bell, D. E. and Padula, L. J. L. (1973). Secure computer systems: Mathematical foundations. Technical Report MTR-2547, The MITRE Corporation, Bedford, MA.

Biba, K. J. (1977). Integrity considerations for secure computer systems. Technical report, The MITRE Corporation.

Bishop, M. (2019). *Computer Security: Art and Science*. Pearson Education Inc., second edition.

Blaze, M., Feigenbaum, J., and Keromytis, A. D. (1998). Trust management for public –key infrastructures. In *Proceedings of the Ninth International Workshop on Services Computing (Lecture Notes in Computer Science 1550)*, pages 59 – 63.

Blaze, M., Feigenbaum, J., and Lacy, J. (1996). Decentralized trust management. In *Proceeding of 1996 IEEE Symposium on Security and Privacy*, pages 164 – 173.

Bonamy, P. (2016). *Maia and Mandos: Tools for Integrity Protection on Arbitrary Files*. PhD thesis, Michigan Technological University.

Clark, D. D. and Wilson, D. R. (1987). A comparison of commercial and military computer security policies. In *IEEE Symposium of Security and Privacy*, pages 184 – 194.

Gambette, D. (1988). Trust: Making and breaking cooperative relations. Technical report, Basil Blackwell Ltd.

Gao, W., Zhang, G., Chen, W., and Li, Y. (2009). A trust model based on subjective logic. In *Proceedings of the Fourth International Conference on Internet Computing for Science and Engineering*, pages 272 – 276.

Jøsang, A. (1999). An algebra for assessing trust in certification chains. In *Proceedings of the Networks and Distributed Systems Security (NDSS'99)*.

Jøsang, A. (2001). A logic for uncertain probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279 — 311.

Jøsang, A. (2002). The consensus operator for combining beliefs. *Artificial Intelligence Journal*, 142(1-2):157 – 170.

Microsoft. Windows integrity mechanism,. https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-azod/75e4ff94-ff5f-43d2-b2e4-4c1429c35261.

Oleshchuk, V. (2012). Trust-enhanced data integrity model. In *2012 IEEE 1st International Symposium on Wireless Systems (IDAACS-SWS)*, pages 109 – 112.

William E. Shotts, J. The linux command line,. https://wiki.lib.sun.ac.za/images/c/ca/TLCL-13.07.pdf.