Architectures for Intrusion Tolerant Database Systems

Peng Liu
School of Information Science and Technology
Pennsylvania State University
University Park, PA 16802
pliu@ist.psu.edu

Abstract

In this paper, we propose four architectures for intrusion-tolerant database systems. While traditional secure database systems rely on prevention controls, an intrusion-tolerant database system can operate through attacks in such a way that the system can continue delivering essential services in the face of attacks. With a focus on attacks by malicious transactions, Architecture I can detect intrusions, and locate and repair the damage caused by the intrusions. Architecture II enhances Architecture I with the ability to isolate attacks so that the database can be immunized from the damage caused by a lot of attacks. Architecture III enhances Architecture I with the ability to dynamically contain the damage in such a way that no damage will leak out during the attack recovery process. Architecture IV enhances Architectures II and III with the ability to adapt the intrusion-tolerance controls to the changing environment so that a stabilized level of trustworthiness can be maintained. Architecture V enhances Architecture IV with the ability to deliver differential, quantitative QoIA services to customers who have subscribed for these services even in the face of attacks.

1 Introduction

The visions of Internet applications (e.g., e-commerce) and pervasive computing not only push computations from a computer into everywhere, but also maximize our dependence on networked computing systems. Quickly increased complexity, openness, inter-connection, and interdependence have made these systems more vulnerable and difficult to protect than ever. The inability of existing security mechanisms to prevent every attack is well embodied in several recent large-scale Internet attacks such as the DDoS attack in Feb. 2000 [44]. These accidents convince the security community that traditional *prevention-centric* security is not enough and the need for *intrusion-tolerant* or *attack-resilient* systems is urgent. Intrusion-tolerant systems, with characteristics quite different from traditional secure systems [12, 3, 18, 40, 30], extend traditional secure systems to be able to *survive* or *operate through* attacks. The focus of intrusion-tolerant systems is the ability to continue delivering essential services in the face of attacks. New mechanisms of attack-resilient systems include but are not limited to intrusion detection, fragmentation, replication, migration, masking, isolation, containment, and recovery.

Being a critical component of almost every mission critical information system, database products are today a multi-billion dollar industry. Database systems motivated 32% of the hardware server volume in 1995 [41], and 39% of the server volume in 2000. Improving the intrusion tolerance of database systems has a direct positive impact on the technology that enables a variety of critical, trusted applications such as e-commerce, air traffic control, credit-card, telecommunication control, and electricity and water supply systems, that our everyday live depends on.

However, existing database security mechanisms are very limited in tolerating intrusions. In particular, authentication and access control cannot prevent all attacks; integrity constraints are weak at prohibiting plausible but incorrect data; concurrency control and recovery mechanisms cannot distinguish legitimate transactions from malicious ones; and automatic replication facilities and active database triggers can even serve to spread the damage.

A Multi-Layer Approach to Database Intrusion Tolerance

Making a database system intrusion tolerant requires in general a *multi-layer* approach, since attacks could come from any of the following layers: hardware, OS, DBMS, and transactions (or applications). A multi-layer approach can be developed along two directions: (1) from scratch or (2) using "off-the-shelf" components.

Along the from-scratch direction, tamper-resistant processing environments [39], and trusted OS or trusted DBMS loaders have been applied to close the door on hardware attacks and OS bugs; certified programs [32, 38] and pro-

tective compiler extensions [8] can be applied to close the door on many DBMS bugs; and signed checksums (and a small amount of tamper-resistant storage to keep the signing key) have been used to detect OS-level data corruption [27]. However, the from-scratch approach is usually not a cost-effective approach, and it cannot be used to tolerate authorized transaction-level intrusions. For example, neither trusted OS nor signed checksums can detect or repair the data corruption caused by a malicious transaction issued by an attacker assuming the identify of an authorized user.

Based on "off-the-shelf" components, OS-level attacks have been addressed by several efforts. In [5], (signed) checksums are smartly used to detect data corruption. In [29] a technique is proposed to detect *storage jamming*, malicious modification of data, using a set of special *detect objects* which are indistinguishable to the jammer from normal objects. Modification of detect objects indicates a storage jamming attack. Although these techniques can effectively tolerant OS-level intrusions, they cannot handle authorized but malicious transactions.

Our Focus and Contributions

In this paper, we focus on transaction-level intrusion-tolerance, which, based on the fact that most attacks are from insiders [6], should be the major threat to database systems; and we propose five architectures for intrusion-tolerant database systems. Although built using "off-the-shelf" components, our frameworks cannot (directly) defend against processor, OS, or DBMS-level attacks, when the lower-level attacks are not so serious and when most attacks are from malicious transactions, our framework can still be very effective. Moreover, existing lower-level intrusion-tolerance mechanisms such as those proposed in [39, 27, 5, 29] can be easily integrated into our frameworks to build a multi-layer, intrusion-tolerant database system.

The remainder of the paper is organized as follows. Section 2 discusses some related work. In Sections 3, 4, 5, 6, and 7, we present five intrusion-tolerant database systems architectures. Section 8 concludes the paper.

2 Related Work

Database security concerns the confidentiality, integrity, and availability of data stored in a database. A broad span of research from authorization [13, 34, 16], to inference control [1], to multilevel secure databases [46, 36], and to multilevel secure transaction processing [4], addresses primarily how to protect the security of a database, especially its confidentiality. Intrusion tolerance, however, is seldom addressed.

One critical step towards intrusion-tolerant database systems is intrusion detection (ID), which has attracted many researchers [26, 31]. The existing methodology of ID can be roughly classed as *anomaly detection* [17, 35, 19, 37] or *misuse detection*[10, 14]. However, current ID research focuses on identifying attacks on OS and computer networks. Although there has been some work on database ID [7, 42], these methods are neither application aware nor at the transaction-level.

The need for intrusion tolerance has been recognized by many researchers in such contexts as *information warfare* [12]. Recently, extensive research has been done in general principles of survivability [18, 45, 11], survivability of networks [30], survivable storage systems [47], survivable application development via middleware [33], persistent objects [28], and survivable document editing systems [43].

Some research has also been done in database intrusion tolerance. In [3], a fault tolerant approach is taken to survive database attacks where (a) several useful survivability phases are suggested, though no concrete mechanisms are proposed for these phases; (b) a color scheme for marking damage (and repair) and a notion of integrity suitable for partially damaged databases are used to develop a mechanism by which databases under attack could still be safely used.

Some of the architectures presented in this paper are directly or indirectly proposed, investigated (using detailed system and algorithm designs), and evaluated (using prototypes) by our previous research. In particular, Architecture I is addressed in [2, 24]; Architecture II is addressed in [22, 20]; and Architecture III is proposed in [21, 23]. However, Architectures IV and V are new. Although in [25] we proposed a rule-based adaptation mechanism for intrusion-tolerant database systems, [25] does not give a comprehensive formal model for adaptive intrusion-tolerant database systems, and such a model is presented by Architecture IV. We include Architectures I, II, and III in this paper because (a) we want to provide a comprehensive view of the fundamental problems in intrusion-tolerant database systems and the corresponding set of promising solutions, and (b) the three architectures build the foundation for Architectures IV and V. It should be noticed that our focus is on architecture level issues and the readers may need to refer to other papers for more design and implementation details.

3 Scheme I

Since the property of database *atomicity* indicates that only committed transactions can really change the database, it is theoretically true that if we can detect every malicious transaction before it commits, then we can roll back the transaction before it causes any damage. However, this "perfect" solution is not practical for two reasons. First, transaction execution is, in general, much quicker than detection, and slowing down transaction execution can cause very serious denial-of-service. For example, the Microsoft SQL Server can execute over 1000 (TPC-C) transactions

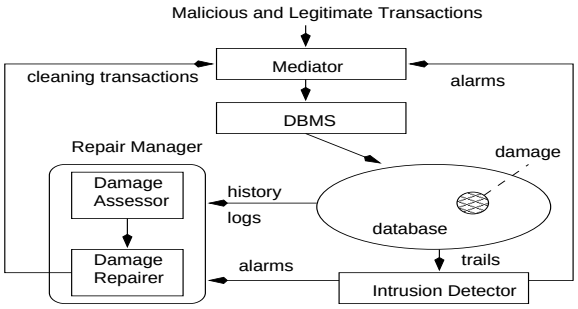

Figure 1. Architecture I

within one second (see *www.oracle.com*), while the average anomaly detection latency is typically in the scale of minutes or seconds. Detection is much slower since: (1) in many cases detection needs human intervention; (2) to reduce false alarms, in many cases a sequence of actions should be analyzed. For example, [19] shows that when using system call trails to identify *sendmail* attacks, synthesizing the anomaly scores of a sequence of system calls (longer than 6) can achieve much better accuracy than based on single system calls.

Second, some authorized but malicious transactions are very difficult to detect. They look and behave just like other legitimate transactions. Anomaly detection based on the semantics of transactions (and the application) may be the only effective way to identify such attacks, however, it is very difficult, if not impossible, for an anomaly detector to have a 100% detection rate with reasonable false alarm rate and detection latency.

Hence, a *practical* goal should be: "After the database is damaged, locate the damaged part and repair it as soon as possible, so that the database can continue being useful in the face of attacks." In other words, we want the database system to operate through attacks.

Architecture I, as shown in Figure 1, combines intrusion detection and attack recovery to achieve this goal. In particular, the *Intrusion Detector* monitors and analyzes the *trails* of database sessions and transactions in a real-time manner to identify malicious transactions as soon as possible. Alarms of malicious transactions, when raised, will be instantly sent to the *Repair Manager*, which will locate the damage caused by the attack and repair the damage. During the whole intrusion detection and attack recovery process, the database continues executing new transactions.

Although there are a lot of anomaly detection algorithms (for host or network based intrusion detection) [17, 35, 19, 37], they usually cannot be directly applied in malicious transaction detection, which faces the following unique challenges:

- Application semantics must be captured and used. For example, for a school salary management application, a $3000 raise is normal, but a $10000 raise is very *abnormal*. Application semantics based intrusion detection is *application aware*. Since different applications can have very different semantics, general application-aware database intrusion detection systems must support dynamic integration of application semantics. Since different anomaly detection algorithms may be good for different application semantics, a general application-aware database intrusion detection system must adapt itself to application semantics.

- Multi-layer intrusion detection is usually necessary for detection accuracy. First, proofs from application layer, session layer, transaction layer, process layer, and system call layer should be *synthesized* to do intrusion detection. Lower level proofs can help identify higher level anomalies. Second, OS-level and transaction-level intrusion detection should be coupled with each other.

We suggest a flexible *cartridge-like* detector to address these challenges. The detector is a cartridge which should be general enough to plug in a variety of (a) anomaly detection algorithms, (b) application semantics extraction algorithms, and (c) application semantics based adaptation policies. The user should be able to prepare some of these algorithms and policies. The detector should provide the interfaces for the user to pick existing and provide new *bullets*, and the detector should not be required to rebuild itself again and again to support each new bullet. (Here each bullet indicates an algorithm or a policy that the detector wants to plug in.) In this way, one detector can be used to meet the intrusion detection needs of multiple applications. Flexibility and expressiveness are the key challenges for developing such a detector. In [15], we have developed a simple cartridge like detector where bullets are supported through DLL modules and a rule based mechanism is used to build the cartridge.

Malicious transactions can seriously corrupt a database through a vulnerability denoted as *damage spreading*. In a database, the results of one transaction can affect the execution of other transactions. When a transaction T_i reads a data object x updated by another transaction T_j, T_i is directly *affected* by T_j. If a third transaction T_k is affected by T_i, but not directly affected by T_j, T_k is indirectly affected by T_j. It is easy to see that when a (relatively old) transaction B_i that updates x is identified as malicious, the damage to x can spread to every object updated by a *good* transaction that is affected by B_i, directly or indirectly. In a word, the read-from dependency among transactions forms the *traces* along which damage spreads.

The job of attack recovery is two-fold: damage assessment and repair. In particular, the job of the *Damage As-*

sessor is to locate each affected good transaction, i.e., the damage spreading traces; and the job of the *Damage Repairer* is to recover the database from the damage caused on the objects updated along the traces. In particular, when an affected transaction T is located, the Damage Repairer builds a specific *cleaning* transaction to *clean* each object updated by T (and not cleaned yet). Cleaning an object is simply done by restoring the value of the object to its latest undamaged version.

Temporarily stopping the database will certainly make the attack recovery job simpler since the damage will no longer spread and the repair can be done backwardly after the assessment is done, that is, we can repair the database by simply undoing the malicious as well as affected transactions in the reverse order of their commit order. However, since many critical database servers need to be 24*7 available and temporarily making the database shut down can be the real goal of the attacker, on-the-fly attack recovery which never stops the database is necessary in many cases.

On-the-fly attack recovery faces several unique challenges. First, we need to do repair forwardly since the assessment process may never stop. Second, cleaned data objects could be re-damaged during attack recovery. Finally, the attack recovery process may never *terminate*. Since as the damaged objects are identified and cleaned new transactions can spread damage if they read a damaged but still unidentified object, so we face two critical questions. (1) Will the attack recovery process terminate? (2) If the attack recovery process terminates, can we detect the termination?

To tackle challenge 1, we must ensure that a later on cleaning transaction will not accidentally damage an object cleaned by a previous cleaning transaction. To tackle challenge 2, we must not mistake a cleaned object as damaged, and we must not mistake a re-damaged object as already cleaned. To tackle challenge 3, our study in [2] shows that when the damage spreading speed is quicker than the repair speed, the repair may never terminate. Otherwise, the repair process will terminate, and under the following three conditions we can ensure that the repair terminates: (1) every malicious transaction is cleaned; (2) every identified damaged object is cleaned; (3) further (assessment) scans will not identify any new damage (if no new attack comes).

From a state-transition angle, the job of attack recovery is to get a *state* of the database, which is determined by the values of the data objects, where (a) no effects of the malicious transactions are there and (b) the work of good transactions should be retained as much as possible. In particular, transactions transform the database from one state to another. Good transactions transform a good database state to another good state, but malicious transactions can transform a good state to a damaged one. Moreover, both malicious and affected (good) transactions can make an already damaged state even worse. We say a database state S_1 is *better* than another one S_2 if S_1 has fewer corrupted objects. The goal of on-the-fly attack recovery is to get the state better and better, although during the repair process new attacks and damage spreading could (temporarily) make the state even worse.

Architecture I has the following properties: (1) It builds itself on top of an "off-the-shelf" DBMS. It does not require the DBMS kernel be changed. It has almost no impact on the performance of the database server except that the *Mediator* can cause some service delay and the cleaning transactions can make the server busier. (2) The intrusion-tolerance processes are all on-the-fly. (3) During attack recovery, the data integrity level can vary from time to time. When the attacks are intense, damage spreading can be very serious, and the integrity level can be dramatically lowered. In this situation, asking the *Mediator* to slow down the execution of new transactions can help *stabilize* the data integrity level, although this can cause some availability loss. This indicates that integrity and availability can be two conflicting goals in intrusion tolerance. (4) More availability loss can be caused when (a) the Intrusion Detector raises false alarms; or (b) a corrupted object is located (It will not be accessable until it is cleaned. Making damaged parts of the database available to new transactions can seriously spread the damage). (5) Inaccuracy of the Intrusion Detector can cause some damage to not be located or repaired. (6) Architecture I is not designed to and cannot handle physical world attack recovery, which usually requires many additional activities. Logically repairing a database does not always indicate that the corresponding physical world damage can be recovered.

To justify the cost-effectiveness of Architecture I, we have implemented a prototype of Architecture I on top of an Oracle database server. Our preliminary testing measurements suggest that when the accuracy of the Intrusion Detector is satisfactory, Architecture I can effectively locate and repair damage on-the-fly with a reasonable impact on (database) performance [24].

4 Scheme II

One problem of Architecture I is that during the *detection latency* of a malicious transaction B, i.e., the duration from the time B commits to the time B is detected, damage can seriously spread. The reason is that during the detection latency many innocent transactions could be executed and affected. For example, if the detection latency is 2 seconds, then Microsoft SQL Server can execute over 2000 transactions during the latency on a single system, and they can access the objects damaged by B freely (since we do not know which objects are damaged by B during the latency).

Quicker intrusion detection can mitigate this problem, however, reducing detection latency without sacrificing the false alarm rate or the detection rate is very difficult, if not impossible. When the detection rate is decreased, more damage is left unrepaired. When the false alarm rate is in-

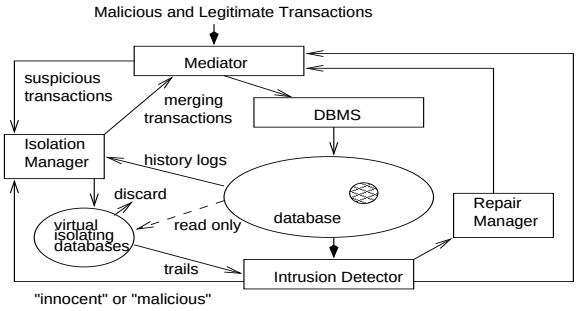

Figure 2. Architecture II

creased, more denial-of-service will be caused. These two outcomes contradict the goal of Architecture I.

Architecture II, as shown in Figure 2, integrates a novel isolation technique to tackle this problem. In particular, first, the Intrusion Detector will raise two levels of alarms: when the (synthesized) anomaly of a transaction (or session) is above Level 1 anomaly threshold TH_m, the transaction is reported malicious; when the anomaly is above Level 2 anomaly threshold TH_s (but below TH_m), the transaction is reported *suspicious*. (The values of TH_m and TH_s are determined primarily based on the statistics about previous attacks). Suspicious transactions should have a significant probability that they are an attack. Second, when a malicious transaction is reported, the system works in the same way as Architecture I. When a suspicious transaction T_s is reported, the Mediator, with the help of the *Isolation Manager*, will redirect T_s (and the following transactions submitted by the user that submits T_s) to a *virtually* separated database environment where the user will be isolated. Later on, if the user is proven malicious, the Isolation Manager will discard the effects of the user; if the user is shown innocent, the Isolation Manager will *merge* the effects of the user back into the main database. In this way, damage spreading can be dramatically reduced without sacrificing the detection rate or losing the availability to good transactions.

We enforce isolation on an user-by-user basis because the transactions submitted by the same user (during the same session) should be able to see the effects of each other. And the framework should be able to isolate multiple users simultaneously. Isolating a group of users within the same virtual database can help tackle collusive attacks, however, a lot of availability can be lost when only some but not all members of the group are malicious. Using a completely replicated database to isolate a user has two drawbacks: (1) it is too expensive; (2) new updates of unisolated users are not visible to isolated users. In Architecture II, we use data *versions* to virtually build isolating databases. In particular, a data object x always has a unique trustworthy version, denoted $x[main]$. And only if x is updated by an isolated user can x have an extra suspicious version. In this way, the total number of suspicious versions will be much less than the number of main versions.

The isolation algorithm has two key parts: (1) how to perform the read and write operations of isolated users (Note that unisolated users can access only the main database); and (2) how to do merging after an isolated user is proven innocent. For part 1, we can enforce *one-way* isolation where isolated users can read main versions if they do not have the corresponding suspicious versions, and all writes of isolated users must be performed on suspicious versions. In this way, the data freshness to isolated users is maximized without harming the main database.

The key challenge in part 2 is the inconsistency between main versions and suspicious versions. If a trustworthy user and an isolated user update the same object x independently, $x[main]$ and the suspicious version will become inconsistent, and one update has to be backed out in order to do consistent merging. In addition, [22] shows that (1) even if they do not update the same object, inconsistency could still be caused; and (2) the merging of the effects of one isolated user could make another still being isolated history invalid. These inconsistencies must be resolved during a merging (e.g., [22] proposes a *precedence-graph* based approach that can identify and resolve all the inconsistencies).

Architecture II has the following set of properties. (1) Isolation is, to large extent, transparent to suspicious users. (2) The extra storage cost for isolation is extremely low. (3) The data consistency is kept before isolation and after merging. (4) During a merge, if there are some inconsistencies, some isolated or unisolated transactions have to be backed out to resolve these inconsistencies. This is the main cost of Architecture II. Fortunately, the simulation study done in [9] shows that the back-out cost is only about 5%. After the inconsistencies are resolved, the merging can be easily done by *forwarding* the left updates of the isolated user to the main database. (5) Architecture II has almost no impact on the performance of the database server except that during each merging process (a) the isolated user cannot execute new transactions; and (b) the main database tables involved in the update forwarding process will be temporarily locked.

We are now implementing an isolation subsystem prototype to further justify the cost-effectiveness of Architecture II [20]. In order to transparently isolate a transaction on top of a commercial single-version DBMS such as Oracle, we need to (a) use extra tables to simulate multiple versions and (b) rewrite the SQL statements involved in this transaction in such a way that the one-way isolation policy can be achieved. Note that query rewriting could cause some service delay to isolated users but not to unisolated users.

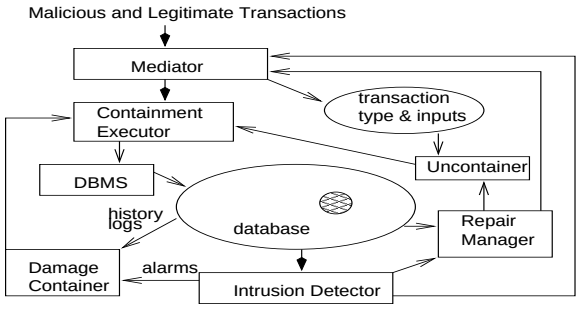

Figure 3. Architecture III

5 Scheme III

Another problem of Architecture I is that its damage containment may not be effective. Architecture I *contains* the damage by disallowing transactions to read the set of data objects that are identified (by the Damage Assessor) as corrupted. This *one-phase* damage containment approach has a serious drawback, that is, it cannot prevent the damage caused on the objects that are corrupted but not yet located from spreading. Assessing the damage caused by a malicious transaction B can take a substantial amount of time, especially when there are a lot of transactions executed during the detection latency of B. During the *assessment latency*, the damage caused during the detection latency can spread to many other objects before being contained.

Architecture III, as shown in Figure 3, integrates a novel multi-phase damage containment technique to tackle this problem. In particular, the damage containment process has one containing phase, which instantly contains the damage that *might* have been caused (or spread) by the intrusion as soon as the intrusion is detected, and one or more later on uncontaining phases to uncontain the objects that are mistakenly contained during the containing phase, and the objects that are cleaned. In Architecture III, the *Damage Container* will enforce the containing phase (as soon as a malicious transaction is reported) by sending some containing instructions to the *Containment Executor*. The *Uncontainer*, with the help from the Damage Assessor, will enforce the uncontaining phases by sending some uncontaining instructions to the Containment Executor. The Containment Executor controls the access of the user transactions to the database according to these instructions.

When a malicious transaction B is detected, the containing phase must ensure that the damage caused directly or indirectly by B will be contained. In addition, the containing phase must be quick enough because otherwise either a lot of damage can leak out during the phase, or substantial availability can be lost. Time stamps can be exploited to achieve this goal. The containing phase can be done by just adding an access control rule to the Containment Executor, which denies access to the set of objects updated during the period of time from the time B commits to the time the containing phase starts. This period of time is called the *containing-time-window*. When the containing phase starts, every active transaction should be aborted because they could spread damage. New transactions can be executed only after the containing phase ends.

It is clear that the containing phase *overcontains* the damage in most cases. Many objects updated within the containing time window can be undamaged. And we must uncontain them as soon as possible to reduce the corresponding availability loss. Accurate uncontainment can be done based on the reports from the Damage Assessor, which could be too slow due to the assessment latency. [21] shows that transaction *types* can be exploited to do much *quicker* uncontainment. In particular, assuming that (a) each transaction T_i belongs to a transaction type $type(T_i)$ and (b) the *profile* for $type(T_i)$ is known, the *read set template* and *write set template* can be extracted from $type(T_i)$'s profile. The templates specify the kind of objects that transactions of $type(T_i)$ can read or write. As a result, the *approximate* read-from dependency among a history of transactions can be quickly captured by identifying the read-from dependency among the types of these transactions. Moreover, the type-based approach can be made more accurate by *materializing* the templates of transactions using their inputs before analyzing the read-from dependency among the types.

Architecture III has the following set of properties. (1) It can ensure that after the containing phase no damage (caused by the malicious transaction) leaks out. (2) As a result, the attack recovery process needs only to repair the damage caused by the transactions that commit during the containing time window, and the termination problem addressed in Architecture I does not exist any longer. (3) One-phase containment and multi-phase containment are the two extremes of the spectrum of damage containment methods. In particular, one-phase containment has maximum damage leakage (so minimum integrity) but maximum availability, while multi-phase containment has zero damage leakage (so maximum integrity) but minimum availability. In the middle of the spectrum, there could be a variety of approximate damage containment methods that allow some damage leakage.

Architectures II and III share the same goal, that is, to reduce the extent of damage spreading, while they take two very different approaches. We are pleased to find that these two architectures are actually complementary to each other and can be easily integrated into one architecture, as illustrated in Figure 4.

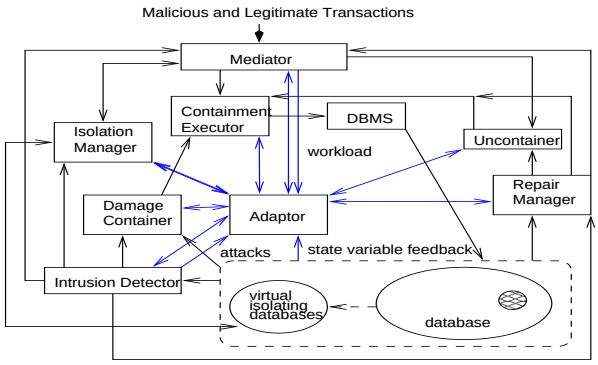

Figure 4. Architecture IV

6 Scheme IV

The intrusion-tolerance components introduced in Architectures I, II, and III can behave in many different ways. At one point of time, the *resilience* or *trustworthiness* of an intrusion-tolerant database system is primarily affected by four factors: (a) the current attacks; (b) the current workload; (c) the current system state; and (d) the current defense *behavior* of the system. It is clear that based on the same system state, attack pattern, and workload, two intrusion-tolerant database systems (of the same Architecture) with different behaviors can yield very different levels of resilience. This suggests that one defense behavior is only good for a limited set of *environments*, which are determined by factors (a), (b), and (c). To achieve the maximum amount of resilience, intrusion tolerant systems must *adapt* their behaviors to the environment.

Architecture IV, as shown in Figure 4, integrates a *reconfiguration* framework to handle this challenge. In particular, an *Adaptor* is deployed to *monitor* the environment changes and *adjust* the behaviors of the intrusion tolerance components in a way such that the adjusted system behavior is more (cost) effective than the old system behavior in the changed environment.

In Architectures I, II, and III, almost every intrusion-tolerance component is reconfigurable and the *behavior* of each such component is *controlled* by a set of a *parameters*. For example, the major control parameters for the Intrusion Detector are TH_m and TH_s. The major control parameter for the *Damage Container* is the amount of allowed damage *leakage*, denoted DL. When $DL = 0$, multi-phase containment is enforced; when there is no restriction on DL, one-phase containment is enforced. The major control parameter for the Mediator is the transaction delay time, denoted DT. When $DT = 0$, transactions are executed in full speed; when DT is not zero, transaction executions are slowed down. At time t, we call the set of control parameters (and the associated values) for an intrusion tolerance component C_i, the *configuration* (vector) of C_i at time t, and the set of the configurations for all the intrusion-tolerant components, the *configuration* of the intrusion-tolerant system at time t. In Architecture IV, each reconfiguration is done by adjusting the system from one configuration to another configuration.

The goal of Architecture IV is to improve the resilience of the system, which has three major aspects: (1) how well the level of data integrity is maintained in the face of attacks; (2) how well the level of data and system availability is maintained in the face of attacks; and (3) how well the level of cost effectiveness is maintained in the face of attacks.

To do optimal reconfiguration, we want to find the best configuration (vector) for each (new) environment. However, this is very difficult, if not impossible, since the *adaptation space* of Architecture IV systems contains an exponential number of configurations. To illustrate, the simplest configuration of an Architecture IV system could be $[TH_m, TH_s, DL, DT]$, then the size of the adaptation space is $domain(TH_m) \times domain(TH_s) \times domain(DL) \times domain(DT)$, which is actually huge. Moreover, we face conflicting reconfiguration criteria, that is, trustworthiness and cost conflict with each other, and integrity and availability conflict with each other. Therefore, we envision the problem of finding the best system configuration under multiple conflicting criteria a NP-hard problem.

Architecture IV focuses on near optimal heuristic adaptation algorithms which can have much less complexity. For example, a data integrity favored heuristic can work as follows: when the level of data integrity, i.e., LI, is below a specific warning threshold I_w, (a) switch the system to multi-phase containment, i.e., let $DL = 0$; (b) slow down the execution of new transactions by $DT = DT + \alpha(I_w - LI)$; and (c) lower the anomaly levels required for alarm raising, that is, $TH_m = TH_m - \beta(I_w - LI)$, and $TH_s = TH_s - \gamma(I_w - LI)$. In this way, we reject and isolate more transactions. Here the values of α, β, and γ are determined based on previous experiences. Note that it is very possible that different (value) combinations of (α, β, γ) are optimal for different environments. Hence it is worthy to have multiple such heuristics with different combinations of (α, β, γ).

It is clear that under different environments different heuristics are the most effective. For example, in some cases integrity favored heuristics can be better, but in some other cases availability favored heuristics can be better. Architecture IV systems should have a mechanism to guide the system to pick the right heuristic (for the current environment). For example, a rule-based mechanism can be used for this purpose.

7 Scheme V

The resilience achieved by Architecture IV is *state-oriented* survivability, that is, the amount of resilience or trustworthiness achieved by Architecture IV is specified, measured, and delivered in terms of the database *state*. For example, at time t, an integrity level of 0.92 achieved by an intrusion-tolerant database system that protects a database of 10,000 data objects can simply mean that 800 objects are corrupted, and an availability level of 0.98 can simply mean that only 200 objects are not accessable. Note than Architecture IV does not *differentiate* between data objects.

Unfortunately, state-oriented, intrusion-tolerant database systems have one serious drawback, that is, they are in general not cost-effective in handling people's intrusion-tolerance requirements in the real world. In the real world, different users usually have different intrusion-tolerance requirements on the shared database system. For example, in a bank, customer Alice could be able to tolerate much less fraud loss on her accounts than Bob on his. In other words, Alice has a much higher integrity level requirement than Bob. In this situation, to satisfy both Alice and Bob, Architecture IV has to achieve (and maintain) the integrity level required by Alice across the whole database, and as a result Architecture IV can waste substantial resources to protect Bob's accounts.

The drawback of state-oriented survivability motivates the idea of *service-oriented* survivability where users' intrusion-tolerant requirements are associated with each (transaction processing) *service*, and the database system's goal is to make sure that the amount of resilience requirement associated with a service is satisfied when the service is delivered. In particular, we call a service associated with a specific level of trustworthiness a Quality of Information Assurance (QoIA) service. And from the viewpoint of users, the goal of a service-oriented, intrusion-tolerant database system is enabling people to get the *QoIA services* that they have subscribed for even in face of attacks. To illustrate, in the above example a QoIA balance inquiry service delivered to Alice could be associated with either one of the following two trustworthiness levels: (1) above 90% accounts involved in this service are not corrupted; (2) for each account involved in this inquiry, the balance reported is at least 90% of the correct balance.

It should be noticed that state-oriented survivability and service-oriented survivability are closely related to each other. Their relationship can be captured by the notions of *state trustworthiness*, which is dependent on the extent to which the data objects can be corrupted or made unavailable, and *service trustworthiness*, which is dependent on the extent to which a service can be distorted by the attacker. If we assume that the DBMS and all transaction codes are trusted, then it is not difficult to see that the QoIA requirements associated with a service can be equivalently *mapped* to a set of state trustworthiness requirements since each service can be modeled as a function of the database state on which the service is executed.

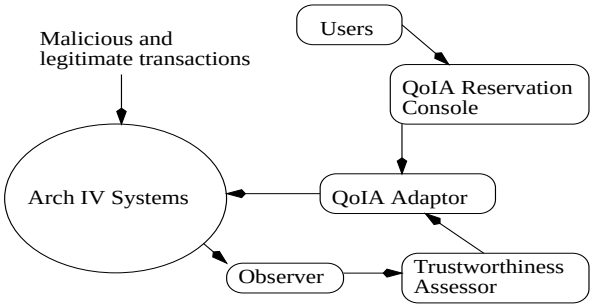

Figure 5. Architecture V

Architecture V, as shown in Figure 5, extends state-oriented, intrusion-tolerant database systems to service-oriented, intrusion-tolerant database systems. In particular, the *QoIA Reservation Console* enables users to subscribe for QoIA services. The *Observer* monitors (and measures) the trustworthiness or healthiness of the database state. The *Trustworthiness Assessor* uses the observed healthiness measurements to *infer* the "real" healthiness of the database state. The *QoIA Adaptor* enhances the Architecture IV Adaptor with the ability to map QoIA requirements associated with services to a set of state trustworthiness requirements and the ability to maintain *differential* state trustworthiness. The adaptation operations performed by the QoIA Adaptor are determined based on the difference between the inferred set of state trustworthiness measurements and the set of state trustworthiness requirements mapped from user QoIA requirements.

To develop an Architecture V system, we face several key challenges. First, although the QoIA requirements associated with a service can be straightforwardly specified based on the results and outputs of the service, delivering a set of QoIA services in a differential way is challenging. Our idea is to indirectly deliver QoIA services through differential state trustworthiness maintenance via the mapping from QoIA requirements to state trustworthiness requirements. Although it is not very difficult to map one service's QoIA requirements to a set of state trustworthiness requirements based on the "function" performed by the service, it could be difficult to resolve the inconsistencies among the set of different state trustworthiness requirements that the set of QoIA services have on a shared data object. Second, how can we maintain differential state trustworthiness? Our idea is to apply different intrusion tolerance controls on different parts of the database. To make this idea feasible, we need to make sure that one set of intrusion-tolerance controls does not influence another

set of intrusion-tolerance controls. Third, how do we ensure that the (mapped) state trustworthiness requirements on a part of the database can be satisfied in the face of attacks? Our idea is through QoIA-aware adaptations where the set of intrusion-tolerance controls enforced on a part of the database can adapt to the changing environment in such a way that the set of state trustworthiness requirements can be satisfied with minimum *cost*. To make this idea feasible, we need to be able to accurately *measure* state trustworthiness. However, this is not an easy job. The measurements observed by the Observer are usually incomplete and could even be misleading due to false negatives, false positives, and detection delays. New techniques are needed to infer the "real" trustworthiness of the database state based on the observed measurements. For example, a statistics based approach could work for this purpose.

8 Conclusion

In this paper, we have presented five intrusion-tolerant database-systems architectures which can be built on top of COTS components. These architectures indicate that: (1) a multi-layer, defense-in-depth approach, as summarized in Figure 6, is usually more cost-effective than having the system's survivability depend on the effectiveness of one or two mechanisms such as intrusion detection; (2) adaptive intrusion-tolerant mechanisms are usually more cost-effective than pre-programmed intrusion tolerant mechanisms; (3) service-oriented, intrusion-tolerant database systems are usually more cost-effective than state-oriented, intrusion-tolerant database systems.

Finally, we would like to mention a couple of exciting future research directions that should be able to further improve the proposed architectures:

- Malicious transactions may be able to be *masked* by a set of partially replicated database servers where each server executes only a group of but not all transactions. The key challenge for such a masking framework should be the tradeoff between security and data consistency.

- It is in general true that the accuracy and latency of the Intrusion Detector can have a big impact on the overall cost-effectiveness of an intrusion-tolerant (database) system. Hence it is very desirable to know how "good" a detector needs to be (in terms of false positive rate, false negative rate, and detection latency) in order to make an intrusion tolerant database system (of Architectures I, II, III, IV, or V) that deploys the detector, cost-effective.

- OS-level and transaction-level intrusion-tolerance mechanisms should be seamlessly integrated to build multi-layer, intrusion-tolerant database systems. This integration requires careful study of the relationships between these two layers of mechanisms. For example, although OS-level data corruptions cannot be detected using transaction-level approaches, transaction-level approaches can be very useful to recover from these corruptions.

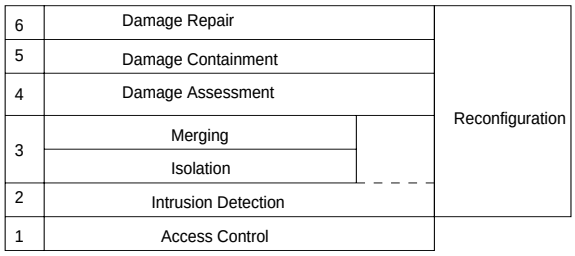

Figure 6. Intrusion Tolerance in Depth

Acknowledgements

This work is supported by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Material Command, USAF, under agreement number F30602-00-2-0575, by DARPA and AFRL, AFMC, USAF, under award number F20602-02-1-0216, by NSF CCR-0233324, and by Department of Energy Early Career PI Award.

References

[1] M. R. Adam. Security-Control Methods for Statistical Database: A Comparative Study. *ACM Computing Surveys*, 21(4), 1989.

[2] P. Ammann, S. Jajodia, and P. Liu. Recovery from malicious transactions. *IEEE Transactions on Knowledge and Data Engineering*, 2002. To appear.

[3] P. Ammann, S. Jajodia, C.D. McCollum, and B.T. Blaustein. Surviving information warfare attacks on databases. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 164–174, Oakland, CA, May 1997.

[4] V. Atluri, S. Jajodia, and B. George. *Multilevel Secure Transaction Processing*. Kluwer Academic Publishers, 1999.

[5] D. Barbara, R. Goel, and S. Jajodia. Using checksums to detect data corruption. In *Proceedings of the 2000 International Conference on Extending Data Base Technology*, Mar 2000.

[6] Carter and Katz. Computer Crime: An Emerging Challenge for Law Enforcement. *FBI Law Enforcement Bulletin*, 1(8), December 1996.

[7] C. Y. Chung, M. Gertz, and K. Levitt. Demids: A misuse detection system for database systems. In *14th IFIP WG11.3 Working Conference on Database and Application Security*, 2000.

[8] C. Cowan, C. Pu, D. Maier, H. Hinton, P. Bakke, S. Beattie, A. Grier, P. Wagle, and Q. Zhang. Stackguard: Automatic adaptive detection and prevention of buffer-overflow attacks. In *Proc. 7th USENIX Security Symposium*, 1998.

[9] S. B. Davidson. Optimism and consistency in partitioned distributed database systems. *ACM Transactions on Database Systems*, 9(3):456–581, September 1984.

[10] T.D. Garvey and T.F. Lunt. Model-based intrusion detection. In *Proceedings of the 14th National Computer Security Conference*, Baltimore, MD, October 1991.

[11] K. Goseva-Popstojanova, F. Wang, R. Wang, G. Feng, K. Vaidyanathan, K. Trivedi, and B. Muthusamy. Characterizing intrusion tolerant systems using a a state transition model. In *Proc. 2001 DARPA Information Survivability Conference (DISCEX)*, June 2001.

[12] R. Graubart, L. Schlipper, and C. McCollum. Defending database management systems against information warfare attacks. Technical report, The MITRE Corporation, 1996.

[13] P. P. Griffiths and B. W. Wade. An Authorization Mechanism for a Relational Database System. *ACM Transactions on Database Systems*, 1(3):242–255, September 1976.

[14] K. Ilgun, R.A. Kemmerer, and P.A. Porras. State transition analysis: A rule-based intrusion detection approach. *IEEE Transactions on Software Engineering*, 21(3):181–199, 1995.

[15] S. Ingsriswang and P. Liu. Aaid: An application aware transaction-level database intrusion detection system. Technical report, Dept. of Information Systems, UMBC, 2001.

[16] S. Jajodia, P. Samarati, V. S. Subrahmanian, and E. Bertino. A unified framework for enforcing multiple access control policies. In *Proceedings of ACM SIGMOD International Conference on Management of Data*, pages 474–485, May 1997.

[17] H. S. Javitz and A. Valdes. The sri ides statistical anomaly detector. In *Proceedings IEEE Computer Society Symposium on Security and Privacy*, Oakland, CA, May 1991.

[18] J. Knight, K. Sullivan, M. Elder, and C. Wang. Survivability architectures: Issues and approaches. In *Proceedings of the 2000 DARPA Information Survivability Conference & Exposition*, pages 157–171, CA, June 2000.

[19] W. Lee and D. Xiang. Information-theoretic measures for anomaly detection. In *Proc. 2001 IEEE Symposium on Security and Privacy*, Oakland, CA, May 2001.

[20] P. Liu. Dais: A real-time data attack isolation system for commercial database applications. In *Proceedings of the 17th Annual Computer Security Applications Conference*, 2001.

[21] P. Liu and S. Jajodia. Multi-phase damage confinement in database systems for intrusion tolerance. In *Proc. 14th IEEE Computer Security Foundations Workshop*, Nova Scotia, Canada, June 2001.

[22] P. Liu, S. Jajodia, and C.D. McCollum. Intrusion confinement by isolation in information systems. *Journal of Computer Security*, 8(4):243–279, 2000.

[23] P. Liu and Y. Wang. The design and implementation of a multiphase database damage confinement system. In *Proceedings of the 2002 IFIP WG 11.3 Working Conference on Data and Application Security*, 2002.

[24] P. Luenam and P. Liu. Odar: An on-the-fly damage assessment and repair system for commercial database applications. In *Proceedings of the 2001 IFIP WG 11.3 Working Conference on Database and Application Security*, 2001.

[25] P. Luenam and P. Liu. The design of an adaptive intrusion tolerant database system. In *Proc. IEEE Workshop on Intrusion Tolerant Systems*, 2002.

[26] T.F. Lunt. A Survey of Intrusion Detection Techniques. *Computers & Security*, 12(4):405–418, June 1993.

[27] U. Maheshwari, R. Vingralek, and W. Shapiro. How to build a trusted database system on untrusted storage. In *Proceedings of 4th Symposium on Operating System Design and Implementation*, San Diego, CA, October 2000.

[28] D. Malkhi, M. Reiter, D. Tulone, and E. Ziskind. Persistent objects in the fleet system. In *Proc. 2001 DARPA Information Survivability Conference (DISCEX)*, June 2001.

[29] J. McDermott and D. Goldschlag. Towards a model of storage jamming. In *Proceedings of the IEEE Computer Security Foundations Workshop*, pages 176–185, Kenmare, Ireland, June 1996.

[30] D. Medhi and D. Tipper. Multi-layered network survivability - models, analysis, architecture, framework and implementation: An overview. In *Proceedings of the 2000 DARPA Information Survivability Conference & Exposition*, pages 173–186, CA, June 2000.

[31] B. Mukherjee, L. T. Heberlein, and K.N. Levitt. Network intrusion detection. *IEEE Network*, pages 26–41, June 1994.

[32] G. C. Necula. Proof-carrying code. In *Proc. 24th ACM Symposium on Principles of Programming Languages*, 1997.

[33] P. P. Pal, J. P. Loyall, R. E. Schantz, and J. A. Zinky. Open implementation toolkit for building survivable applications. In *Proc. 2000 DARPA Information Survivability Conference (DISCEX)*, June 2000.

[34] F. Rabitti, E. Bertino, W. Kim, and D. Woelk. A model of authorization for next-generation database systems. *ACM Transactions on Database Systems*, 16(1):88–131, 1994.

[35] D. Samfat and R. Molva. Idamn: An intrusion detection architecture for mibile networks. *IEEE Journal of Selected Areas in Communications*, 15(7):1373–1380, 1997.

[36] R. Sandhu and F. Chen. The multilevel relational (mlr) data model. *ACM Transactions on Information and Systems Security*, 1(1), 1998.

[37] S. Sekar, M. Bendre, and P. Bollineni. A fast automaton-based method for detecting anomalous program behaviors. In *Proc. 2001 IEEE Symposium on Security and Privacy*, Oakland, CA, May 2001.

[38] Z. Shao, B. Saha, and V. Trifonov. A type system for certified binaries. In *Proc. 29th ACM Symposium on Principles of Programming Languages*, 2002.

[39] S. Smith, E. Palmer, and S. Weingart. Using a high-performance, programmable secure coprocessor. In *Proc. International Conference on Financial Cryptography*, Anguilla, British West Indies, 1998.

[40] V. Stavridou. Intrusion tolerant software architectures. In *Proceedings of the 2001 DARPA Information Survivability Conference & Exposition*, CA, June 2001.

[41] P. Stenstrom and et al. Trends in shared memory multiprocessing. *IEEE Computer*, (12):44–50, December 1997.

[42] S. Stolfo, D. Fan, and W. Lee. Credit card fraud detection using meta-learning: Issues and initial results. In *Proc. AAAI Workshop on AI Approaches to Fraud Detection and Risk Management*, 1997.

[43] M. Tallis and R. Balzer. Document integrity through mediated interfaces. In *Proc. 2001 DARPA Information Survivability Conference (DISCEX)*, June 2001.

[44] C. Taylor. Behind the hack attack. *Time*, (2):45–47, February 2000.

[45] F. Webber, P. P. Pal, R. E. Schantz, and J. P. Loyall. Defense-enabled applications. In *Proc. 2001 DARPA Information Survivability Conference (DISCEX)*, June 2001.

[46] M. Winslett, K. Smith, and X. Qian. Formal query languages for secure relational databases. *ACM Transactions on Database Systems*, 19(4):626–662, 1994.

[47] J. J. Wylie, M. W. Bigrigg, J. D. Strunk, G. R. Ganger, H. Kiliccote, and P. K. Khosla. Survivable information storage systems. *IEEE Computer*, (8):61–68, August 2000.