

On Achieving Fast Damage Appraisal in case of Cyber Attacks

Chandana Lala and Brajendra Panda¹
Computer Science Department
University of North Dakota
Grand Forks, ND 58202, USA
Email: {lala, panda}@cs.und.edu

Purpose

Invasion of information systems through electronic media has become rampant with the outburst of Internet technologies and their applications. An information attack brings a system to unstable and inconsistent state [1], [3]. In such cases, the system under attack cannot differentiate an attacker from a legitimate user. Once the attacking transaction modifies some data in the database and commits, its effect becomes permanent and available to other users. Then, the damage can spread to other parts of the database through legitimate users as they update fresh data after reading damaged data [2], [4]. The damage will propagate through the database with delay in recovery. Hence it is absolutely imperative to perform immediate damage assessment and recovery, as soon as an attack is detected, to stop the cascading effect of the damage. The objective of this research is to develop accurate and fast damage assessment techniques to recover the database in real-time in the event of an intrusion.

Method

In traditional recovery methods, after an intrusion detection mechanism identifies an attack, the log of the affected database would be scanned starting from the attacking transaction until the end. Then, the effect of the attacker and all affected transactions would be undone. Next, the affected transactions would be re-executed. This requires significant amount of work. Our idea is to obtain the dependency relationships among transactions when they execute. This relationship is then stored in an auxiliary structure, which can be memory resident. In the event of an intrusion, this structure could be used, as opposed to the log, to determine affected transactions. We have developed four algorithms based on this approach. We have used transaction commit sequences instead of transaction IDs in order to facilitate the search. The first algorithm utilizes a dependency list to store transaction dependency graphs. The second one uses a precursor list. For every transaction T_i , this list stores the IDs all other transactions whose updates were read by T_i . While the third algorithm uses a byte format to store dependencies, the fourth one uses a bit-vector for every transaction to capture the information. Advantages and disadvantages of each method have been examined.

Results

When used, any of the developed algorithms can precisely determine the transactions affected by the attacker. During recovery, only the tampered portions of the log are accessed and a major part of it is skipped. This helps in reducing the recovery process dramatically. Currently, we are working on performance analysis of the model using simulation. Initial results indicate that each of the four structures expedite the recovery process considerably.

New or Breakthrough Aspects of Work

The existing damage assessment and recovery algorithms require reading a major portion of the database log in order to identify the affected transactions. This will involve significant page I/Os, thus, delaying the process. Our algorithms use auxiliary structures, which capture the transaction dependency relationships as transactions commit. These structures are extremely compact compared to the log. During damage assessment and recovery, each developed algorithm processes corresponding structure and determines affected transactions. Then the portions of the log involving these transactions are read and affected data items are recovered. Thus, our model significantly reduces damage assessment and recovery process, as a result, diminishing the denial-of-service attack time.

Conclusions

In case of an information attack, a malicious transaction can corrupt parts of databases and, without immediate attention, the damage can spread through the system very quickly. Hence quick and efficient damage appraisal and recovery method is vital for survivability of information systems [5]. This research offers four different algorithms to perform fast

¹ This work was partially funded by US AFOSR grant F49620-99-1-0235

damage assessment and recovery on databases. When an attack is detected, these algorithms identify affected transactions and their respective starting addresses in the log quickly and correctly without accessing the log. We are developing a simulation analysis of the model and the initial results show considerable improvement over traditional method.

References

- [1] P. Ammann, S. Jajodia, C. D. McCollum, and B. Blaustein, "Surviving Information Warfare attacks on Databases", In Proceedings of the 1997 IEEE Symposium on Security and Privacy, pp.164-174, Oakland, CA, May 1997.
- [2] R. Graubart, L. Schlipper, and C. McCollum, "Defending Database Management Systems Against Information Warfare Attacks", Technical Report, The MITRE Corporation, 1996.
- [3] S. Jajodia, C. D. McCollum, and P. Amman, "Trusted Recovery", In Communications of the ACM, Vol. 42, No. 7, p. 71-75, July 1999.
- [4] P. Liu, P. Ammann, and S. Jajodia, "Rewriting Histories: Recovering from Malicious Transactions," Distributed and Parallel Databases, Vol. 8, No. 1, p. 7-40, January 2000.
- [5] B. Panda and J. Giordano, "Defensive Information Warfare", In Communications of the ACM, Vol. 42, No. 7, p. 31-32, July 1999.