

Architecture and Components for Data Management Security: NRL Perspective

C. E. Landwehr
J. N. Froscher

*Naval Research Laboratory
Center for High Assurance Computer Systems
Code 5542
Washington, D.C. 20375-5337*

The DoD urgently needs system architectures that will permit the use of commercial off-the-shelf technology (COTS), including database management systems, without making the system as a whole vulnerable to COTS shortcomings and defects. Practical architectures to meet this need will, we believe, assign security functions that require high assurance to separate, simple components, and use COTS components for the purposes for which they are developed.

The commercial sector has introduced mechanisms for access control, integrity, authentication, privacy, and non-repudiation that are designed to counter the threats to which commercial systems are subject. If the DoD can organize its systems so that there are environments where the commercial threat model holds, it may very well be appropriate to rely on the commercial mechanisms to counter those threats. However, DoD still needs to provide users with different clearances access to classified data and services at different classification levels (i.e., DoD users need multilevel secure (MLS) systems). DoD may also wish to detect and recover from the effects of subtle attacks on commercial databases, to provide cleared users with access to unclassified databases that resists traffic analysis, and to provide more secure and convenient means for maintaining user authentication to the DBMS. The balance of this position paper describes technology now being developed and/or demonstrated at NRL that provides the means to address DoD's unique security policy requirements for data management using small scale, inexpensive components in conjunction with COTS database technology.

Preventing Insecure Information Flow

Industry does not in general recognize information flow vulnerabilities as a serious threat and has pursued MLS solutions only with support from government. Commercial database vendors lack the incentive to enforce DoD information flow policies. Custom databases that enforce these policies will be created and maintained only if there is continuing DoD support, and such non-standard systems are likely to

- be incompatible with commercial applications,
- lack features users desire that competing commercial databases will have, and
- cost substantially more than commercial databases.

But if DoD develops its systems based on architectures that assign the enforcement of its specialized security policies to separate, simple, high assurance, low cost components,

operating together with commercial DBMSs, it creates the opportunity for MLS database service based on COTS databases.

A few researchers in the MLS community have developed prototype client and server solutions that allow the use of commercially available components with simple, high assurance devices to control the flow of information across classification boundaries. These efforts provide high assurance, effective separation of classified information. These approaches exploit the physical separation inherent to distributed computing and the secure transmission of low information to high systems. They also rely on the INTELINK infrastructure that DoD and the Intelligence Community have installed over the past few years. It is important to note that both client and server technology must be able to provide high users access to lower classified information.

Two client solutions becoming available are the Starlight Interactive Link [1] and the COSPO Switched Workstation. The Starlight Interactive Link (IL), developed by the Australian Defence Science and Technology Organisation and currently being commercialized, allows the user of an X-Windows workstation to switch the keyboard and mouse output dynamically between an X-server on a High network and a (physically distinct) X-server on a Low network. The workstation's monitor is connected to the High X-server only. The results of any keystrokes and mouse-clicks sent to the low X-server are transmitted to the High X-server through a one-way flow device, called a data diode, and displayed in a Low window on the user's display. Starlight allows cut-and-paste from Low to High, and the Low window can be iconified while work on the High server is in progress.

A slightly different approach is taken by the COSPO Switched Workstation, which has been approved for use between Secret and Unclassified networks and can be used with any PC-based workstation. The Switched Workstation uses a commercially available component called a server switch¹ [2]. In this scheme, the user switches both the monitor and the keyboard/mouse between a server on a High network and a server on a Low network by means of a specific sequence of keystrokes recognized by the server switch. A separate one-way flow device (e.g., a one-way fiber optic link, with supporting software running on the low and high sides) allows the transmission of Low data to High systems; supporting software on the Low side can check data integrity and screen data for viruses. This approach does not provide a windowing capability.

In both of these client-side approaches, when Low data are propagated to High systems, the High copy of the Low data is protected at the High level while a Low copy remains on the Low server. When users wish to release the Low data to Low users or systems, both Starlight and the Switched Workstation allow them to release the data at its correct classification without any need for a "downgrade" operation.

¹ Server switches were developed to meet a commercial need to share a single monitor and keyboard among several servers when, for example, an administrator wants to use only a single keyboard and monitor to control several machines (e.g. servers or routers). The commercial technology is applied here to switch a keyboard and monitor between systems at different security levels.

A similar approach to the server side has been demonstrated in several relational database prototypes developed by the SINTRA (Secure INformation Through the Replicated Architecture) project [3]. Physically separate COTS DBMS products at different security levels contain information at their level and below, so that users at a given session level have access to data at that level and below. When users update the database at their level, the update must be securely and consistently replicated to the higher databases. With the commercial availability of asynchronous replication servers, the only MLS critical feature required for this approach is a one-way communication component that ensures secure information flow as well as reliability, fairness, availability, and performance. The NRL Pump is an example of a device that balances these requirements [4]. The SINTRA approach can be applied to many data management paradigms but depends on the commercial availability of an asynchronous communication protocol. A MLS CORBA is being prototyped using this approach and the Starlight IL.

These approaches can provide high assurance that DoD information flow policy is enforced while requiring little or no trusted software. They can provide MLS database service for DoD users without requiring an MLS database because they minimize the need to share physical resources among users and processes at different security levels.

Detecting Subtle Attacks on Databases

Commercial DBMS systems provide a variety of integrity controls that are designed to detect inconsistencies that may arise through error or accident. Some database applications also include checks designed to prevent or detect fraud. None of these mechanisms, however, are designed to detect or resist subtle attacks that aim simply to degrade the operation of the DBMS rather than to enrich the perpetrator or completely disable the database. Such attacks, termed *jamming* or *data spoofing*, could nevertheless pose a significant threat to DoD operations. NRL researchers have recently developed and prototyped approaches that can work with COTS DBMSs to detect, and, in some cases, recover from such attacks [5].

Protecting Database Accesses Against Traffic Analysis

As the problem statement for this panel observes, a user with a laptop computer and a cellular phone can now access and gather information from databases around the world. Unfortunately, unless such a user takes substantial precautions, it is also possible for every router in the path between the user and the database to eavesdrop on that communication. And even if the user chooses to encrypt the contents of the communication with the database, the header information, which reveals the source and destination of the traffic the user sends, can be read by every router. Defense users who wish to gather open source intelligence may well prefer not to reveal their tracks to all and sundry.

Onion routing [6] provides a way to protect such database accesses (and Internet connections generally) against traffic analysis. This approach calls for the user's web browser to access the Internet via an onion routing proxy, which operates at the application level, perhaps on the firewall at the user's site. On receiving a request to connect to a particular web site, the onion routing proxy chooses a route through a network of onion-routers, creates a layered data structure (an onion) that corresponds to that route, and sends the onion to the next onion router in the path. The layers of the onion are multiply-encrypted, and each onion router in the path can strip off only one layer of encryption, so that it can learn the identity only of its immediate neighbors in the path. Once the connection is set up, data sent over the path is similarly wrapped in layers of encryption, so that it appears in the clear only after it exits the final onion router.

The application-layer onion routers can be hosted by firewalls, as applications running on independent workstations, or, potentially, on a user's own workstation. To take advantage of onion routing, a user need only configure her web browser to use an onion-router for an HTTP proxy (onion-routing proxies will also be available for other protocols as well). No modifications to browsers or other proxy-aware applications are required.

Continuously Authenticating DBMS Users

The current state of the practice for user authentication in virtually all DoD DBMSs is the user password and ID. Although this system, properly implemented and administered, can provide reasonable authentication, it is prone to a number of well known problems -- poorly chosen or

shared passwords, passwords compromised through sniffers, passwords written on yellow stickies attached to the workstation, and so on. Further, once the user is authenticated, it is usually easy for the user to walk away from the workstation without logging out, leaving his identity to be borrowed by anyone with physical access to the workstation.

In the past year, NRL researchers have prototyped a system that combines an inexpensive token based on automotive remote key entry technology and a detector that can be plugged to any PC or Macintosh to provide more effective and continuous user authentication [7]. Again, the idea is to offload the security critical function onto a separate, simple, low cost but high assurance device. A user announces his arrival by pressing a button on the token, which initiates continuing communication between the token and a detector. The workstation's keyboard and monitor are plugged into the detector; if the detector loses contact with the user's token, it breaks the connections between the processor and the keyboard/mouse and monitor. This system can potentially be extended to deal with a variety of operational environments and may be adapted to provide the user's password to the processor

The Future

Where once were distinct, monolithic databases that could be independently managed and secured, we now find an increasingly interconnected and interdependent system of systems. Re-engineering of business processes typically causes workflows among various parts of an organization to become more highly automated and more tightly integrated. For DoD to maintain the security of its databases as these changes occur requires a system architecture that places DoD-specific security requirements on individual components that can be built inexpensively and with high assurance. Further, the architecture must accommodate the inter-operation of frequently changing COTS components without introducing unacceptable security risks.

The database client solutions and the server solution described in this paper can be combined to produce a comprehensive architecture for distributed computing that supports MLS computing services [8]. This approach allows users access to information at their level and below and allows them to originate data at its correct classification level. Technology alone cannot eliminate the need for downgrades; DoD users must adopt a different approach to accessing classified information. However, there is still a need for automatic downgrade of sanitized information. High assurance downgraders can provide this function but should be disconnected from the global confederation when the downgrades are executed. Encryption protects information in transmission across public networks. By establishing environments within the overall system to which commercial threat models can reasonably be applied, DoD can exploit commercial approaches to privacy, integrity, authentication, non-repudiation, and access control. For example, commercially available, B1 DBMS products can reasonably be used to control access to compartmented information within a classification level in the way earlier environmental guidelines recommended.

With this change in its philosophy of protection, DoD can realize much stronger protection of classified information, more widespread secure use of COTS products, reduced cost, a

migration path for the secure use of legacy resources, secure integration of new technology, and promotion of information sharing in its system of heterogeneous, autonomous systems.

REFERENCES

- [1] Anderson, M., North, C., Griffin, R., Yesberg, J., and Yiu, K., "Starlight: Interactive Link," *Proc. Twelfth Annual Computer Security Applications Conf.*, San Diego, CA, Dec., 1996.
- [2] Examples of such server switches can be found in the catalogs of Black Box Corp., Pittsburgh, PA., available at URL <http://www.blackbox.com/> or Computer Ware OnLine, available at URL <http://www.cwol.com/>
No product endorsement is implied by these references.
- [3] Froscher, J. N., M. H. Kang, J. P. McDermott, O. Costich, and C. E. Landwehr, "A Practical Approach to High Assurance Multilevel Secure Computing Service," *Proc. Tenth Annual Computer Security Applications Conference*, Orlando, FL, Dec., 1994.
- [4] Kang, M. H., Moskowitz, I. S. and Lee, D. C., "A Network Pump" *IEEE Transactions on Software Engineering*, Vol. 22, No. 5, May, 1996.
- [5] McDermott, J. P. and J. N. Froscher, "Replay and Replication Defenses Against Storage Jamming," *Proc. 20th National Information Systems Security Conf.*, 1997 (this proceedings).
- [6] Syverson, P. F., D. M. Goldschlag, and M. G. Reed, "Anonymous Connections and Onion Routing," *Proc. 1997 IEEE Symp. on Security and Privacy*, pp. 44-54.
- [7] Landwehr, C. E., "Protecting Unattended Computers Without Software," submitted for publication, June, 1997.
- [8] Kang, M. H., Froscher, J. N., and Moskowitz, I. S., "A Framework for MLS Interoperability," *Proceedings of IEEE HASE Workshop*, Oct. 1996.