

LARSTAN'S  
**THE BLACK  
BOOK ON**

**CORPORATE  
SECURITY**

Multi-Level Security:  
Your Key to Data Safety

**JIM PORELL**

[6]

# MULTI-LEVEL SECURITY: YOUR KEY TO DATA SAFETY

*Multi-level security, with its capacity for sharing resources across several security compartments, provides more efficient and, ultimately more secure data workflow.*

**"HISTORY DOES NOT  
TEACH THAT BETTER  
TECHNOLOGY NECESSARILY  
LEADS TO VICTORY. RATHER,  
VICTORY GOES TO  
THE COMMANDER  
WHO USES TECH-  
NOLOGY BETTER."**

- Office of the Chief of Naval Ops.

by **JIM PORELL**

**D**uring the Cold War, it was called "fail-safe," meaning a multi-level nuclear warfare system structured with enough redundancies to make the probability of failure extremely remote. During this new era of information warfare, think of multi-level security (MLS) as a fail-safe data system.

In the broadest definition, MLS consists of sharing resources across separate security compartments.

## **A MULTI-LEVEL SECURITY ENVIRONMENT HAS PROTECTION LEVELS BUILT IN TO STOP SOMEONE FROM DECLASSIFYING, MISCLASSIFYING, OR RECLASSIFYING DATA.**

These resources might represent data, applications and networks. A security compartment can take different forms, whether as hierarchical levels of security, or as specific company projects. It is also defined by determining who has access to what information. By sharing, a lot of redundancy is eliminated in an MLS, as compared to the commonly deployed alternative, the multiple security levels (MSL).

In an MSL environment, each department only has whatever it needs to function. It may have its own network, its own applications and its own data. These multiple departments, or multiple communities, each with its own computing infrastructure, get quite costly, especially when data and applications have to be replicated.

By contrast, the concept behind MLS is to share data and, ultimately, applications to reduce operational complexity. In addition, the timeliness of data is ensured since the time lag required to replicate data and sanitize it for different departments is eliminated. The policies to backup, protect and secure the data are also consistently applied in this shared environment.

As I will explain, MLS results in data operations that are streamlined, consolidated, more efficient, less costly and considerably more secure.

### **COMPARING MLS TO MSL**

In an MSL environment, an analyst has a PC on their desktop. This person will have a distinct network that's identified only to run data that is specific for that PC or for the project that the person is working on. This network is then connected into a server, which may have access to data defined for that department. This same analyst might work across three separate departments, and might have three separate computers in their office, with three separate networks all going to three independent servers, with each of those servers having its own specific data. There is no mechanism to share. That analyst can't copy, paste or move data from one computer to another.

---

MLS is about finding opportunities for sharing. In the last example, some data that comes in on a project might be relevant to each of the three areas identified, as well as other areas. It may arrive initially for one area, but then could be shared with another. To make this information accessible to them, it can be copied, or the analyst can be granted access to this other area. However, the data's classification needs to be protected. If it came in secret, it can't be declassified. A multi-level security environment has protection levels built-in to stop someone from declassifying, misclassifying or reclassifying data.

## MULTI-LEVEL SECURITY ACCESS ELEMENTS

Two security elements are central to a multi-level security environment:

- Discretionary access control (DAC), under which an employee is a member of a group, and that group has access to resources. In turn, the employees get access to the resources.
- Mandatory access control (MAC), under which an employee has to be explicitly defined to have access to those resources. It implies certain process rules and functions as an enforcement mechanism. It dictates who is allowed to see this data, who is allowed to reclassify this data and, in some cases, who is allowed to know this data exists.

An example of MAC: an individual is in the insurance department of a company that does banking, as well. It would be good for that individual to know how much is in a customer's account, so insurance policies could be tailored to them. If a customer has \$1 million in the bank, the insurance policy would differ from someone with \$10,000 in his account. In some cases, that sharing of personally identifiable information is against the law. Federal regulations governing information use, such as the Health Insurance Portability and Accountability Act (HIPAA) and Sarbanes-Oxley, say this information should not be shared. The insurance depart-

**Insider Notes:** A major opportunity for an MLS commercial deployment is in the privacy area to enforce regulations like Sarbanes-Oxley, HIPAA and the California Privacy Acts. The compartmentalizing of business processes that require mandatory access control will become more prevalent.

## **MLS RECOGNIZES THAT CORPORATE INFORMATION IS JUST ALL ONE GIANT WORKFLOW THAT CAN BE MANAGED BY SHARING DATA, APPLICATIONS AND OPERATIONS.**

ment cannot see any of the personally identifiable information from the banking department. Enforcing access to that kind of information is one of the commercial applications of multi-level security.

### **THE COMPONENTS OF MLS**

The initial impetus for MLS has been in the government sector, in particular the intelligence community. However, there is a growing need and desire to leverage this same capability in the commercial space.

Three system components that should be optimized for an MLS system:

- the database server
- the applications server
- the presentation device

Each component can have elements of multi-level security applied to IT. The state-of-the-art is that there are several existing database servers that can provide a level of compartmentalization. There are currently no general-purpose, off-the-shelf application servers with this capability commercially available. Some specific ones are built to government specifications. There are some efforts now being undertaken to build thin client terminals that will have a level of compartmentalization in them. The focus for the industry today is in the database server itself. Other than customized specific applications, there are no credible end-to-end MLS solutions in the marketplace today.

### **COMMERCIALIZING MLS**

There's definitely a desire to develop this MLS capability for the commercial marketplace, but some developmental work on hardware is still required. The presentation server is the next component being developed and optimized for MLS. Application servers are also in development, but it will be several years before any robust MLS-aware application servers come to market.

---

The storage capability required for MLS is already deployed, since a database server is dependent on its storage device. Therefore, tape devices and some of the Redundant Array of Independent Drives (RAID) storage devices can run as compartmentalized data stores. MLS capabilities, such as erase-on-scratch, were added to the hardware devices, since it is no longer acceptable to just remove a file from a directory on a storage device. It must be ensured that nobody can take that storage device offline and find residual data. Business processes also need to be in place to look at removable media and online media and determine its life span, so that nobody can go off and pick up the remnants of some data because it has been partially written over.

A major opportunity for an MLS commercial deployment is in the privacy area to enforce regulations like Sarbanes-Oxley, HIPAA and the California Privacy Acts. The compartmentalizing of business processes that require mandatory access control will become more prevalent. The world is a distributed computing environment, and no one computer can meet all of a businesses' needs. The mainframe can't do it because it's blind and deaf. It does not own the human computer interface, so data has to be moved to a device with that capability.

A large mainframe might have all of a company's transaction files in one place. However, these files are accessed and replicated to run in a business intelligence application on UNIX servers. Each time this data is replicated, the preservation of policies protecting that data, like privacy, security levels, or access control, are put in jeopardy. It's very easy to copy data; it's extremely difficult to copy the policy and to have the same policy being implemented on a completely different type of server. A critical success factor to facilitate the sharing of data is to have consistent policies that manage resources, and, more importantly, that can be utilized for audit, compliance and risk assessment against these resources.

**Insider Notes:** A critical success factor to facilitate the sharing of data is to have consistent policies that manage resources, and, more importantly, that can be utilized for audit, compliance and risk assessment against these resources.

This is a fundamental point where companies need to start changing their whole perception about how they do computing. Currently, many companies opt to use the cheapest and smallest computers, so they copy, cluster and replicate. This process makes the data accessible to a broad range of servers. However, the key issue is whether the company is consistent in maintaining privacy and management policies with each replicated instance of that data. Companies had an audited server in the first instance where the data was generated, but did they have a similar audit on each subsequent instance of that data?

When companies start looking at the new government mandates and realize that they are very complex and require new software and new workflows to facilitate policy management, they will question why these policy decisions are all independent and why they are not incorporated as part of the corporate workflow.

MLS also means getting away from the stovepipe mentality that dictates separate and distinct computers for customer relationship management applications, for human resource applications and for point of sale transactions. MLS recognizes that corporate information is just all one giant workflow that can be managed by sharing data, applications and operations. By doing this, costs can be reduced dramatically, along with risk exposure and the need for continual, complex compliance analysis.

Companies are initially handling these newly mandated security regulations as a matter of just going through the old audit checklist. In reality, compliance can be automated if the volume of replications is minimized. MLS is one method to minimize replications and to consolidate servers. It also starts to incorporate and simplify policy management along with server consolidation.

Forty years ago, someone would flowchart an application to determine which tasks it was supposed to accomplish. But now, with business process integration, data is moving all over the enterprise through these silos of computing. This is the new business workflow. The question is whether this movement is orchestrated or simply being passed to the next decision-making unit — from the point of sale application to the customer relationship management to target marketing, or web-based management. If


---

## DATA SNAPSHOT

This isolated approach towards managing userids separately can lead to errors in attempting to correlate security across multiple systems.

*Download the complete research study for free at [www.blackbooksecurity.com/research](http://www.blackbooksecurity.com/research)  
Source: 2005 Larstan / Reed Infosecurity Survey*

**Are user IDs managed at an enterprise level (across servers) or against each individual server?**



this entire workflow is considered, perhaps the better decision is to move some of the application, not the data. Flowcharting the data can be a tremendous exercise in reducing complexity. Eliminating data movement can save time, which saves money. When looked at from the policy perspective and the cost of managing policy against the data, there are even more opportunities for tremendous operational savings.

MLS takes what might have been operational issues that the security administrator does and goes back to the fundamentals of business architecture and business workflows, and ultimately makes the business organization more efficient.

Below, I define the traditional security paradigm and discuss its inadequacies, and then present how MLS differs from the traditional approach. I discuss the opportunity for MLS, how it would work and its implications for executive, IT and security management.

## THE TRADITIONAL SECURITY PARADIGM

The traditional data security paradigm is based on the three As: authentication, authorization and access control. Audit capabilities may also be included. This paradigm focuses on a discrete user or group of users and defines

**Insider Notes:** The traditional data security paradigm is based on the three As: authentication, authorization and access control.

## **MLS IS JUST BECOMING GENERALLY AVAILABLE. IT PROVIDES COMPANIES WITH FLEXIBILITY AND PRESENTS A CHANCE TO CHANGE SECURITY MANAGEMENT POLICY. IT SIMPLIFIES OPERATIONS.**

their access to a fixed resource such as data, a printer or a particular program. However, patterns of use change as the role of the user changes. For example, a systems programmer could have access to some system data set during a test environment to facilitate change management. When you get into a production environment where you have to manage change a lot more tightly, this person may still have authorized access to that data set, but the company does not want to facilitate change management there. That is the evolution of access control.

Access control has existed in the government for some time, but now, because of Sarbanes-Oxley, HIPAA and the like, it is also evolving into the commercial space. How it gets implemented is yet to be determined. This is the opportunity to promote some of the concepts and philosophies of MLS as a way to manage and achieve this evolution.

An example of how MLS works: a person in a banking role may have a particular access to data, but if he is working in the same business in a stock or insurance role, he may not be authorized to that same level of information for that specific task. Part of MLS is to identify what role a user is playing or what compartment they are working within for a particular task environment, and then to enforce the security policies for that role. The user could have access to all kinds of information when looking at all the different roles that they could have in their job. But in a given task, they are restricted access to certain information, either by law, by corporate policy or by customer initiated request, and they are not supposed to make decisions based upon this restricted information. Therefore, that information should not be accessible to them during that function.

The key differences between these paradigms are the roles that people play versus the access control that any single individual might have, and defining more policies based on role and upon the life cycle of the data or application.

---

### How does multi-level security change the way companies address the security paradigm?

MLS is just becoming generally available. It provides companies with flexibility and presents a chance to change security management policy. It simplifies operations. New data security regulations offer another opportunity, since there is no programmatic way to enforce these laws. Compliance becomes a paper exercise, an audit trail, a post-processing exercise, or a lot of activity to set up stovepipes of information. What multi-level security does is present a method to share data and systems and facilitate working together without affecting regulatory restrictions.

### THE MULTI-LEVEL SECURITY OPPORTUNITY

Nobody is currently using multi-level security commercially. There are multiple situations for its effective implementation. Multi-level security is focused on roles. In the government intelligence community, it deals with various compartments, considers the task that each individual needs to execute and applies its level of access accordingly.

Commercially, an individual might have access to a broad wealth of data, information and applications within their company. But when they are in a specific role or executing a specific function, then they might (by law, policy or request) only be allowed to do certain tasks.

For example, in the healthcare industry, a claims administrator might have a wealth of knowledge of personally identifiable information about a particular client. This insurance institution might also be selling other forms of insurance or annuities and mutual funds. That personally identifiable information should not be leveraged for the sales of the mutual funds. If they are selling stock, they are not supposed to have all the information

**Insider Notes:** A MLS implementation requires a combination of hardware and software optimized for this type of sharing. The manner in which a server enforces MLS may be proprietary. However, it is then the role of the network to ensure interoperability between two environments that may implement MLS differently.

## **OUTSOURCERS ARE EXTREMELY RIPE FOR AN MLS ENVIRONMENT. LEVERAGING MLS, OUTSOURCERS ARE MORE EFFICIENT BECAUSE THEY HOST MULTIPLE CLIENTS IN THE SAME COMPUTING INFRASTRUCTURE AND ENSURE THAT ONE CLIENT CAN'T SEE THE DATA OF ANOTHER CLIENT.**

that the claims administrator can access. These businesses will restrict an individual to one task only, be it banking, insurance or selling stock.

With MLS, a person can identify themselves to a task or compartment. If they are executing a stock action, they can only use the information that is legally accessible to them. Policy is being enforced based on the business role versus individual user. That individual may have access to personally identifiable information, but only in the banking role, not the insurance or stock selling role. MLS can facilitate this solution. This may allow some companies to reduce their infrastructure and their operational costs by replacing separate, compartmentalized systems with people who can handle multiple business roles.

### **THE OPPORTUNITY IN THE EXTENDED ENTERPRISE AND OUTSOURCING**

The extended enterprise depends on business-to-business communications. All of these interactions start with sharing data, but at some point, content is being shared, which means sharing the application. This results in a person from one company running on a computer of another company. This is similar for an outsourcer.

Outsourcers are extremely ripe for an MLS environment. Leveraging MLS, outsourcers are more efficient because they host multiple clients in the same computing infrastructure and ensure that one client can't see the data of another client. In turn, the outsourcer is saving dramatically in their operational infrastructure. Instead of having a computer system or a collection of servers for each customer, they are sharing that infrastructure across multiple customers.

---

In business-to-business communications, lots of data gets transmitted. An example is the local phone companies that actually do all the metering for long distance calls. At the end of the month, they separate, parse and burn a lot of computing time to distribute that information to the long distance line carriers who, in turn, receive it, and process the bills using their own computing infrastructure. If the long distance carriers and the local carriers could share some of this computing infrastructure, each of them would probably save a considerable amount of computer processing time.

In another example, a company in the content management business, such as eBay, may be able to leverage this infrastructure to conduct private auctions only for people who sign up for it. Others would have no knowledge of, or access to, this private auction. However, the database infrastructure could be shared, saving operational expense for the content hosting service.

## MULTI-LEVEL SECURITY — GETTING IT TO WORK

MLS implementation requires a combination of hardware and software optimized for this type of sharing. The manner in which a server enforces MLS may be proprietary. However, it is then the role of the network to ensure interoperability between two environments that may implement MLS differently. Part of what MLS does is reduce the heterogeneous nature of a system, but still employs virtual private networks, IP security and data encryption over the wire. Those are the security standards that facilitate interoperability. On the local systems, how MLS may be achieved is an implementation detail, and typically, proprietary to that type of server.

On the operations side, synergy is important. This is where standards will emerge. For example, the radio buttons of a management interface, pushed to develop or implement a MLS environment, ultimately requires some

**Insider Notes:** The core of how MLS works is in the mapping process. Compartments are defined by correlating privacy statements or sensitive data to the various functional roles that exist within a company. These policies are then made operational by identifying information compartments that are accessible to employees as they assume various roles within their function.

## EXECUTIVE MANAGEMENT CAN VIEW MLS AS A METHOD TO BREAK DOWN THE ACCESS STOVEPIPES AND REDUCE COSTS BY ENABLING ONE PERSON TO LEGALLY ASSUME SEVERAL ROLES WITHIN THE COMPANY.

consistency. What's behind that radio button can be anything, as long as the operations are consistent. But that is more of a visionary statement, not an implementation detail, and certainly not what is currently the state-of-the-art. As systems become heterogeneous, more consistency is required in operational models to reduce complexity. Here are the implications and benefits of MLS across the business structure:

- **Executive Management.** Executive management can view MLS as a method to break down the access stovepipes and reduce costs by enabling one person to legally assume several roles within the company. This is accomplished by eliminating or reducing data movement, evolving individual roles and providing the computing infrastructure to facilitate a true business role based management philosophy. This should also save operational costs.
  - **IT Management.** MLS will implement the differences in access demanded by a production versus a test environment. Systems programmers in a test role can do anything they want. They can reboot the system and make changes dynamically just to get the product working. Individuals may also have the same responsibilities in the production environment, but not the same freedom or flexibility to do changes. MLS will enforce those roles and that policy.
  - **Security Management.** Security management will be responsible for administering these roles. This will demonstrate that security management can have a large impact on the company by developing this business role-based paradigm that will eventually lead to company savings. Therefore, how security management charts this course and helps identify infrastructure changes necessary for an MLS implementation will help to establish a much more resilient, economically feasible and more cost-competitive operational environment.
-

Now, I'll describe how an MLS system can be implemented.

## UNDERSTANDING THE WORKFLOW AND FUNCTIONAL ROLE

The key part of implementation is in understanding the business workflows within a company and the roles and functions that employees have within it. It is a difficult task to take those business roles and map them into the access control, authentication, authorization and audit policies necessary to define the compartments required by MLS. To accomplish this, a company really needs to focus on the end-to-end workflows of its business processes.

## MAPPING LEVEL OF ACCESS

Once these workflows are understood, a company can map out the data and information to be made accessible to those employees functioning within specific business roles. It can then define the level of access control for sensitive data and information that must be maintained and managed when an employee is in that business role. This is how a company maps security policy into the actual workflow of its infrastructure.

For example, if a person is a bank teller, the company needs to understand what their function is in detail and what data and information access they need to have to do their job on a day-to-day basis. This gets more complex when an employee has a position in selling insurance, as well. If part of his function is to handle client requests, he may have access to personally identifiable information.

However, when that same person is handling unsolicited sales requests, then he can't have access to all of this information, while in that role. Within an MLS structure, the role that an employee is performing is critically important in determining how much information is made available to them at that given time. This is where MLS begins to provide compartmentalization.

**Insider Notes:** A company must identify what personally identifiable information and other sensitive data it needs to protect from general access. It must precisely define when it can use this data to target markets and when it cannot.

## **PUTTING IT ALL TOGETHER, A COMPANY NEEDS TO DEVELOP A ROAD MAP THAT DEFINES AN ACTION PLAN FOR IMPLEMENTATION. IT THEN NEEDS TO BUILD A MATRIX OF SENSITIVE DATA AND PERSONALLY IDENTIFIABLE INFORMATION DISTINCT TO THAT COMPANY.**

Mapping roles to information access is a complex task and may require new tools to get it done right. Often, when a company maps privacy policy and sensitive data access to a corporate function, it finds that it has multiple policies that may affect the same role, or policies that may not let a business role function adequately. This mapping process allows a company to review its security policies in detail and determine how it affects the company's workflow.

The next step is to turn these access control policies into actionable items that can evolve into security compartments. This entails defining the security policies and then mapping specific business roles against them that detail the degree of access to the types of sensitive information that each role should have. This ensures that the rules and responsibilities are clearly understood, and enforced companywide. The goal is to make access control of sensitive information trivial to audit, so compliance with the new regulations are easy to demonstrate and difficult to fault.

The core of how MLS works is in the mapping process. Compartments are defined by correlating privacy statements or sensitive data to the various functional roles that exist within a company. These policies are then made operational by identifying information compartments that are accessible to employees as they assume various roles within their function.

### **THE ACCESS/COMPANY ROLE MATRIX**

Research is now being conducted to utilize employee prose and natural language parsing to identify the various access needs of business roles, such as the needs of a bank teller, a client service representative, a problem management person, an insurance sales person, etc. This automatically

---

**DATA SNAPSHOT**

Understanding who has a need-to-know for particular data simplifies the deployment of a multi-level secure environment.

*Download the complete research study for free at [www.blackbooksecurity.com/research](http://www.blackbooksecurity.com/research)  
Source: 2005 Larstan / Reed Infosecurity Survey*

**Is employee access to data managed on a need-to-know basis?**

**YES: 73%**

**NO: 27%**



helps determine the level of access each of these business roles may have to sensitive data.

In addition, a company needs to identify what personally identifiable information and other sensitive data it needs to protect from general access. It must precisely define when it can use this data to target markets and when it cannot. The company also needs to know what its policy is when a customer checks the box saying that the company can use this information for target marketing, as well as the policy when the company said that it would never sell its list.

The result of these decisions can be put into a matrix that models a company's privacy policy and defines what specific business roles can access and can take actionable efforts against this personally identifiable information and sensitive data. The first step is to determine what the personally identifiable information is (or other sensitive information that a company does not want made available to other business roles or another company), and second, determine how the computing infrastructure can enforce that privacy strategy.

**Insider Notes:** Within an MLS environment, where mandatory access control is being enforced, it can be set up so that an unauthorized user is not even aware that other resources exist. This eliminates the fishing that someone might have to perform to gain access to that data.

## **SENIOR MANAGEMENT MUST BUY IN COMPLETELY BECAUSE MULTI-LEVEL SECURITY IS A WORKFLOW, NOT A STOVEPIPE DECISION. IT'S NOT A SERVER DISCUSSION OR EVEN AN INFORMATION TECHNOLOGY DISCUSSION; IT'S AN ENTERPRISE DISCUSSION.**

To a large extent, MLS and privacy are related because the whole point of this implementation is to hide certain types of information from certain individuals within specific business roles that don't have a need to know. It is basically dealing with some form of need-to-know while also facilitating data and application sharing and reducing operational infrastructure.

### **THE FINAL STEPS**

Putting it all together, a company needs to develop a road map that defines an action plan for implementation. It then needs to build a matrix of sensitive data and personally identifiable information distinct to that company. The company then applies the security policies that define each compartment and details the business roles and information access contained in each, and determines what data and information can be shared, and what cannot be shared. Finally, the company must determine who else may have access. If it is a business partner that has access, then it must be decided what sub-roles within that business partner's company has access to this information, and at what level.

Senior management must buy-in completely because MLS is a workflow, not a stovepipe decision. It's not a server discussion or even an IT discussion; it's an enterprise discussion. It's about how a server implementation of MLS adds value to the enterprise and all the other systems that are connected to it. It's neither a technology nor an individual business unit decision; it's an enterprise-wide decision.

### **DIFFERENCES WITH NON-MLS ENVIRONMENTS**

In a non-MLS environment, discretionary access control, permissions and access control lists are applied against individual users or groups. An employee that is part of a group is allowed to access certain data.

---

However, in these situations, there is no consideration for networking issues or for network enabled applications. The focus is on a discrete resource or object, like a file or a database. MLS differs in that a company has the opportunity to apply access control to the network and to the applications.

In some respects, mandatory access control is implemented in MLS. Another important difference is that an employee may know that a resource exists under the non-multi-level secure environment, but not know what the content of this particular data field or data-set might be, and he doesn't have access to it. He might then try to gain access by asking someone else to provide him with it. Within an MLS environment, where mandatory access control is being enforced, it can be set up so that an unauthorized user is not even aware that other resources exist. This eliminates the fishing that someone might have to perform to gain access to that data.

If the view of data provided for a business role doesn't have any excess information, a user might be aware that some personally identifiable information exists, but unaware of the depth and breadth of that information. Maybe there are eight pieces of data and the user may be cognizant of only one and not aware that there were seven other major data fields that were hidden. This provides a greater level of security, greater isolation and compartmentalization of tasks on behalf of the company.

## SETTING UP MLS - ACCESS CONTROL DECISIONS

Setting up the MLS environment and facilitating mandatory access controls is the first and most difficult step. There are a lot of knobs that need or can be turned to describe access policies. Companies have a tremendous amount of flexibility as to how granular an MLS implementation is required. On one hand, a company could decide that it's just user and key data that they will compartmentalize. They could get down to a field level within a database.

**Insider Notes:** The problem among the mainframe, UNIX and Windows environments is they are typically separate security domains. To simplify the operations and enforce consistent policies, the security policies should be merged.



## CASE STUDY: SECURITY IS A MANY-LAYERED THING

**The Problem:** A financial services company typically encompasses multiple business units. A small company will often have personal banking, insurance and stock brokerage operations; larger companies may have even more.

A variety of laws prohibit companies from leveraging all the information of one business unit (such as personally identifiable information collected within, and for the banking operation) to target market a consumer for other business units (such as the insurance and stock brokerage units).

To meet the requirements of these regulations, the finance company needs to compartmentalize or isolate the applications and data within each business unit to prevent the emergence of conflicts. This will require the company to install redundant systems and storage capacity to facilitate this isolation. In addition, it adds an entire level of complexity for change management activities and opens additional opportunities for error with the required replication of certain data and the need for accurate information on all systems.

### A WORKFLOW FOR THIS FINANCIAL SERVICES COMPANY

A customer enrolls in one of the business units. Information is gathered about that individual to open an account. This enters the customer into the company's Customer Relationship Management (CRM) system. An initial transaction is executed on behalf of that customer, such as depositing money into an account, executing a stock trade or acquiring insurance. This transaction typically occurs on a mainframe system.

Now the other business units within the company are interested in acquiring this consumer's business as well. The operational data store (ODS) for the transaction system will be "mined"; data will be extracted, filtered and sent to a data warehouse or decision support system (DSS). These functions typically reside in a UNIX systems environment.

From there, applications will be run against the "allowable" information in that data warehouse. A target set of "pre-approved" consumers will

---

be defined and a mailing or call will be placed to the consumer in order to interest them in new business elements. Should the consumer accept this new offer, this cycle repeats itself, beginning with updates to the CRM system.

Each time the data moves between these systems, it is important to ensure that the privacy policies are enforced. Audits need to be conducted to ensure that the policies have been enforced and then correlated across each of the systems to ensure there are no irregularities. This post-processing operation can be complex, cumbersome and error prone, especially if it is not kept properly in synch across each of the application execution environments, which span across somewhat independent business units.

## **HOW MULTI-LEVEL SECURITY ADDRESSES THIS PROBLEM**

Multi-level security can be leveraged to alleviate much of this complexity. The data for each of the business units can be shared in a common database. This will require a cross-company effort to ensure that the proper database schema is identified to capture each business units' unique information. Compartments can then be identified at both the row and field level within these databases.

Categories can be set up to identify collections of customers for each individual business unit or customers with activities in multiple business units. For example, in this company, security labels could be established such as BANK, INSURANCE, STOCK, BANKINS, INSSTOCK, BANKSTOCK, BIS. This will create a hierarchy of labels that would be generated for a consumer within the CRM application. The label would be modified, based on which business units the consumer belongs. From that point on, each businesses' applications could operate against the common ODS, but only operate against their specific customers.

For example, a bank reconciliation application might run with a security label of ALLBANK, which includes BANK, BANKINS, BANKSTOCK and BIS. The ODS running MLS would never access any non-banking customers resident within the server and would not impact any summary counts or averages. This would help reduce some database

programming by avoiding additional “WHERE” clauses that avoid or ignore customers.

In the non-MLS environment, each business unit may have had their own decision support system for targeting new opportunities. In the MLS environment, the same decision support system may be shared across the business units. Leveraging applications that specify a security label with their processing provides a mechanism to ensure that specific consumers, identified by their database row security label or secure views against field level information, do not get targeted inappropriately. More importantly, business unit actions can be audited since system options are available to record both failures and successes against specific data within the ODS or DSS.

An MLS system will enable a company with multiple business units that require separate security postures to use the same database and the same applications without breaching security protocol. This will result in a simplified configuration without redundant systems that will provide fully audit-friendly security for sensitive and regulated information.

A company could also decide that only an internal network will support a certain level of information, and from an open network, such as the Internet, the volume of information made accessible will be reduced for security reasons. It can decide that only specific applications can access the data or that certain data can only be viewed within a specific procedure. It is through an access control program that these rules will be decided. In addition, a company can decide to put flashes on printed output, such as confidential, classified, or privacy-act-like statements that state that this information needs to be protected against copying or further dissemination.

## **COMPANY DECISIONS**

The company first needs to determine its critical information processing infrastructure and the critical assets that it needs to protect. The evolution in applying MLS to this infrastructure is in understanding what are critical assets, what are critical data, the applications that access that data and the network that it gets accessed over. You must determine the level of protection necessary to meet a particular government requirement or to enforce a

---

privacy policy. Critical assets can be on stored media, the database itself and anywhere there are files. It may be a field or a row within a relational database. It can be a removable media, tape devices, optical storage carousels or printed forms. Companies make a privacy commitment to their consumers and they need to put this into practice.

## **FLOWCHARTING**

The simplest method to determine access to critical data is to chart the data flow. From that data flow a company can determine who has access to the various elements of that data all through its lifecycle. This really comes down to a lifecycle management structure. With input, process and output visualizations provided by a flowchart or other lifecycle tooling, a company can look at:

- how the data gets from one point to another
- how that transfer is secured
- who had access to that transfer
- what programs accessed that data at each point
- the persons who can access those programs at each point

These capabilities allow a company to begin to understand where it needs to setup its access control and permission lists. But even more importantly, it helps to identify company groupings so mass customization of security can occur. If a company can reduce the number of data moves and shares of information across applications, then it will reduce the complexity of operations.

Most of this implementation is in execution or just in a normal workflow. How the company wants to provide an audit, conduct compliance or complete analysis reports, to determine that it has properly protected its assets, will determine the level of granularity and the completeness of the process. If this is done on behalf of a government regulation, it will be more detailed. Companies would apply similar granularity on behalf of customers.

However, for internal use, the implementation may be less detailed as the company may be willing to take more risk. As a company has greater accountability to either a government entity or its own consumers, it will

## **THE COMPANY FIRST NEEDS TO DETERMINE ITS CRITICAL INFORMATION PROCESSING INFRASTRUCTURE AND THE CRITICAL ASSETS THAT IT NEEDS TO PROTECT.**

want to use rigor in how it implements this audit and compliance management process.

### **RESTRUCTURING THE EXISTING WORKFLOW**


The lifecycle of company data probably entails several replicated data moves and several different applications manipulating new images of this same data in what can be termed a stovepipe approach of operation. It might have a mainframe setup for point of sale and ATM type of transaction processing, a UNIX server setup for some type of decision support and data warehouse environments, and a Windows system setup for human resources applications.

Each one of these application environments is independently managed, but the reality is that the same data is flowing through all those elements. Therefore, if they are independently managing and securing them, they will have difficulty in guaranteeing the security and management policies against the data throughout its life cycle. These companies will have to change some of their management approaches to look at those separate application processes as a single, consistent business workflow and manage the security consistently across it.

The problem among the mainframe, UNIX and Windows environments is they are typically separate security domains. To simplify the operations and enforce consistent policies, the security policies should be merged. The first issue is whether the company can merge any of the data. This can reduce the number of data moves to simplify policy management operations. A data merger would mean that the teams that were working independently before on different aspects of this business workflow would now have to work a lot more closely together.

It may also change some of the infrastructure. It is critical for the company to recognize that these systems will have to be peers in terms of working

---

**DATA SNAPSHOT****How often is failed access to data audited?**

Like searching for shoplifters in the retail industry, if no one ever looks at the videotapes or closed circuit TVs, a business can't stop or inhibit losses. Audit is an important element of preventing and inhibiting security loss.

*Download the complete research study for free at [www.blackbooksecurity.com/research](http://www.blackbooksecurity.com/research)  
Source: 2005 Larstan / Reed Infosecurity Survey*

together, because the company's workflow is distributed and will remain so as long as the company has applications servers on Windows and UNIX systems and the data on a mainframe. There might be a hierarchy in terms of how controls are implemented, but they need to have peer services and consistency in management. It might be different implementations based on server specific needs, but applying consistent security concepts across the platform is a critical success factor.

Not only do systems need to be a data server, but they should also be data clients. A system might have data in a UNIX or Windows environment that a mainframe application wants to have access to. Or a system might be pushing data out from the mainframe using a client file protocol to a Windows or UNIX server just to simplify the workflow.

Where the mainframe has direct access to another system, via a file transfer operation, one system will be a data client and the other will be a data server. The systems must communicate, and there has to be some degree of credential sharing between those systems. One thing that the mainframe can do is function as an enterprise authentication server. This will allow a

**FORTY YEARS AGO, SOMEONE WOULD FLOWCHART AN APPLICATION TO DETERMINE WHICH TASKS IT WAS SUPPOSED TO ACCOMPLISH. BUT NOW, WITH BUSINESS PROCESS INTEGRATION, DATA IS MOVING ALL OVER THE ENTERPRISE THROUGH THESE SILOS OF COMPUTING. THIS IS THE NEW BUSINESS WORKFLOW.**

company to authenticate this business workflow in a single place with a common user identification structure that is then used as the access control mechanism on each independent server. By centralizing user registration enrollment and authentication, the company will start to pull together security controls across the enterprise. That will assist in reducing operational complexity.

### **HOW TO MAKE THIS WORK?**

One of the simplest steps for a company to consider is how it provides registration, identification, and authentication of users, who defines the roles, and how they are defined. Progressing from there, the company starts to consider how to share some of the data within common user credentials. If the business starts looking at centralizing some aspects of authentication across the business workflow, it should find that the sharing of data or facilitating the sharing of data will become operationally simpler. This in turn yields a reduction in storage management operations against that data.

Multi-level security is focused on the sharing of compartmentalized data. It is intended to manage and enforce common security policies against that managed data. It will certainly take some additional effort to understand what resources can be shared across multiple communities. However, once that activity is completed, a single shared instance of data will reduce the total cost of operations against that data, proportionally to the number of communities that share the data.

---

These savings will be earned by:

- streamlining security operations
- consolidating storage management tasks (such as backup and recall)
- reducing the number of physical server and storage devices
- lessening the complexity of audit and compliance post-processing activities
- shortening the time to make information accessible to additional communities, because now the business will not have to wait for a replication or copy of the data

The fact is, multi-level security facilitates data sharing that can yield significant operational savings to the customer across many fronts.



*Jim Porell is a Senior Technical Staff Member and Chief Architect for IBM's mainframe software. He chairs the zSeries Software Design council that brings together all IBM technology utilized in the deployment of customer solutions utilizing mainframe technology. Jim led the architecture for IBM's Multi-Level Security technology for z/OS, including the database server, DB2. These became generally available in March 2004, though the base technology dates back to 1987.*

*Jim focuses on Security and Business Resilient solutions for customers. He has been consulting to government and commercial customers on secure computing infrastructure for 10 years. He has participated in several government efforts ensuring the protection of critical computing infrastructure. He can be reached at 845-435-6593 or [jporell@us.ibm.com](mailto:jporell@us.ibm.com).*